

TeamQuest® Performance Software Statistics Reference Manual

September 2014

Release 11.2 PF 20140922

TQ-40023.4

Copyright © 2014 TeamQuest Corporation. All Rights Reserved.

The names, places and/or events used in this publication are purely fictitious and are not intended to correspond to any real individual, group, company or event. Any similarity or likeness to any real individual, company or event is purely coincidental and unintentional.

NO WARRANTIES OF ANY NATURE ARE EXTENDED BY THE DOCUMENT. Any product and related material disclosed herein are only furnished pursuant and subject to the terms and conditions of a license agreement. The only warranties made, remedies given, and liability accepted by TeamQuest, if any, with respect to the products described in this document are set forth in such license agreement. TeamQuest cannot accept any financial or other responsibility that may be the result of your use of the information in this document or software material, including direct, indirect, special, or consequential damages.

You should be very careful to ensure that the use of this information and/or software material complies with the laws, rules, and regulations of the jurisdictions with respect to which it is used.

The information contained herein is subject to change without notice. Revisions may be issued to advise of such changes and/or additions.

U.S. Government Rights. All documents, product and related material provided to the U.S. Government are provided and delivered subject to the commercial license rights and restrictions described in the governing license agreement. All rights not expressly granted therein are reserved.

TeamQuest, the TeamQuest logo, TeamQuest Alert, TeamQuest Analyzer, TeamQuest Baseline, TeamQuest CMIS, TeamQuest Harvest, TeamQuest IT Service Analyzer, TeamQuest IT Service Reporter, TeamQuest Manager, TeamQuest Model, TeamQuest Predictor, TeamQuest Online, TeamQuest Surveyor, TeamQuest View, and Performance Surveyor are trademarks or registered trademarks of TeamQuest Corporation in the US and/or other countries.
Microsoft and Windows are registered trademarks of Microsoft Corporation in the US and/or other countries.
UNIX is a registered trademark of The Open Group in the US and/or other countries.
All other trademarks and service marks are the property of their respective owners.

About This Manual

TeamQuest Performance Software is a suite of four integrated products that help organizations optimize IT services while minimizing infrastructure costs and mitigating risks. TeamQuest Manager®, a component of the TeamQuest CMIS® product, is the required software that provides the services for the TeamQuest Performance Software suite. When TeamQuest Manager is installed on a host system, it supplies the capabilities to collect, store, manage, and administer your performance data.

Purpose

This manual is to be used as a supplement to the *TeamQuest Performance Software Administration Guide*.

Scope

This manual provides descriptions of the statistics collected by the collection agents of TeamQuest Manager®. This manual also contains derived statistic descriptions for use with TeamQuest Alert® and TeamQuest View®.

Audience

This manual is intended for those who are interested in performance analysis and capacity management of computer systems. TeamQuest Performance Software addresses the needs of the performance analyst and the capacity planner. The performance analyst, who is typically interested in detailed data, can use TeamQuest Manager to collect system data at a fine granularity.

The capacity planner can use TeamQuest Manager to collect data across long periods of time. Marketing representatives and technical managers who are interested in the capacity management of computer systems can also use TeamQuest Manager.

Prerequisites

This manual assumes that you have basic knowledge of the operating system environment. (For example, you can log into the system, edit files, run commands, and so on.)

How to Use This Manual

This manual is to be used with the TeamQuest Performance Software applications that report or use the statistics collected by the TeamQuest Manager collection agents.

Organization

This manual is organized as follows:

Section 1. Introduction

This section provides background information about the parameter hierarchy, table data, and performance data collectors.

Section 2. AutoPredict Statistics

This section provides descriptions of the statistics collected by the AutoPredict Agent.

Section 3. DB2 Universal Database (UDB) Server

This section provides descriptions of the statistics collected by the DB2 UDB Agent.

Section 4. Hewlett-Packard HP-UX Systems

This section provides descriptions of the statistics for Hewlett-Packard HP-UX systems.

Section 5. Hyper-V Statistics

This section provides descriptions of the statistics collected by the Hyper-V Agent.

Section 6. IBM AIX Systems

This section provides descriptions of the statistics for IBM AIX systems.

Section 7. KVM Systems

This section provides descriptions of the statistics collected by the Libvirt Agent.

Section 8. Linux Systems

This section provides descriptions of the statistics for Linux systems.

Section 9. Microsoft Windows Systems

This section provides descriptions of the statistics collected by the System Activity Agent and the Process-Workload Agent. It also provides a description of the derived statistics in the performance database.

Section 10. Network Applications

This section provides descriptions of the statistics collected by the Network Application Agent.

Section 11. Network Devices

This section provides descriptions of the statistics collected by the Network Device Agent.

Section 12. Oracle Database Server

This section provides descriptions of the statistics collected by the Oracle Data Agent and the Oracle Alarm Agent.

Section 13. Oracle Solaris Systems

This section provides descriptions of the statistics for Oracle Solaris systems.

Section 14. PostgreSQL Database Agent

This section provides descriptions of the statistics collected by the PostgreSQL Database Agent for PostgreSQL databases.

Section 15. Sybase ASE Server

This section provides descriptions of the statistics collected by the Sybase ASE Agent.

Section 16. System Alarm Statistics

This section provides descriptions of the statistics recorded in the System Alarm table by the Alarm Service.

Section 17. VMware Systems

This section provides descriptions of the statistics collected by the VMware Infrastructure Agent.

Section 18. Web Server

This section provides descriptions of the statistics collected by the Web Server Agent.

Related Product Information

The following related documents may be helpful to you when using TeamQuest Performance Software. Use the version that corresponds to the level of software in use at your site.

TeamQuest Administration Console User Guide (TQ-40321)

This guide tells you how to use the TeamQuest Administration Console. The TeamQuest Administration Console provides the capabilities to administer a large number of systems running TeamQuest Manager from one central location. This guide includes deployment strategies and recommendations, introductory information, and overall workflow descriptions. It includes information about the overall product and user interface features. It also describes the tasks necessary for all types of users to implement the TeamQuest Administration Console.

TeamQuest Analyzer User Guide (TQ-40242)

This guide tells you how to use TeamQuest Analyzer. TeamQuest Analyzer provides a network accessible, browser-based user interface for detailed data reporting and analysis of your enterprise.

TeamQuest CMIS Installation Guide (TQ-50015)

This guide provides installation instructions for the product components of TeamQuest CMIS. TeamQuest CMIS includes TeamQuest Manager, the PostgreSQL server, and the TeamQuest Administration Console.

TeamQuest Performance Software Administration Guide (TQ-40020)

This guide provides information for using TeamQuest Manager to administer your TeamQuest performance databases. It tells you how to use the TeamQuest Manager interface to configure the collection agents that collect and store data in the performance database. It also provides information for configuring the TeamQuest system and service agents, maintaining the various TeamQuest policies, and performing the day-to-day activities involved when using TeamQuest Manager.

TeamQuest Performance Software Administration Reference Manual (TQ-40022)

This manual contains conceptual information about TeamQuest Manager and provides information to help you customize TeamQuest Manager for your site.

TeamQuest Performance Software Command Line Interfaces Reference Manual (TQ-40024)

This manual provides the formats and descriptions of the command line interfaces available with the TeamQuest Performance Software.

TeamQuest Performance Software Enterprise Solutions Administration Guide (TQ-40212)

This guide tells you how to use TeamQuest Manager and its associated collection and service agents in an enterprise-wide environment. It also provides information about deploying the TeamQuest Performance Software products on the systems in your enterprise, configuring data collection to store data from multiple systems into a common TeamQuest enterprise database or an open database, and using the TeamQuest Update Server.

TeamQuest Performance Software Installation and Configuration Guide for KVM Systems (TQ-17015)

This guide describes how to install TeamQuest Manager for use with Kernel-based Virtual Machine (KVM). It also provides all the configuration information required to collect configuration and performance data for KVM hosts and virtual machines using TeamQuest Manager.

TeamQuest Performance Software Installation and Configuration Guide for VMware vSphere (TQ-18015)

This guide describes how to install TeamQuest Manager for use with VMware vSphere. It also provides all the configuration information required to collect configuration and performance data for VMware vCenter servers, VMware hosts, and virtual machines using TeamQuest Manager.

TeamQuest View Reports Reference Manual (TQ-16028)

This manual contains descriptions of the predefined reports that come with TeamQuest View.

TeamQuest View User Guide (TQ-01401)

This guide describes how to use TeamQuest View. TeamQuest View is a graphical user interface that allows you to perform detailed data reporting and analysis on your PC or workstation. You use TeamQuest View to view reports from data collected by the TeamQuest collection agents.

How to Access Portable Document Format (PDF) Files

All TeamQuest product documentation is available as PDF files. PDF files can either be downloaded from the TeamQuest website or they can be viewed from your product DVD.

You must have the Adobe Reader software from www.adobe.com installed on your workstation to view the PDF files.

To download product documentation from the website:

1. Access the TeamQuest website at www.teamquest.com.
2. Click the Support menu, then Customer Area.
3. Log in to the Customer Area.
4. Click Download products, patches and documents.
5. Click Product Documentation for the desired release level and product platform.
6. Click document you want to download.
7. Open or save the PDF file.

To access product documentation from your product DVD:

1. Load the DVD in your disk drive.
The disk contents dialog box is displayed.
2. Open the welcome.pdf file.

The welcome.pdf file can be used to navigate to the desired PDF file.

You may wish to copy the welcome.pdf file and the pdfdocs folder to your server so multiple users have convenient access to the product documentation.

Note: *When copying the welcome file and the pdfdocs directory, you must maintain the established directory structure. If the directory structure is not maintained, the navigation links within the welcome.pdf file will not open the documents located in the pdfdocs directory.*

Notation Conventions

In this manual, the following conventions apply:

- *TQDIR* indicates the path of the directory where the TeamQuest software is installed. Whenever you see *TQDIR*, substitute the path of the directory at your site.
- *TQDATADIR* indicates the path of the directory where the performance database is installed. Whenever you see this text, substitute the path of the directory at your site.
- Program names are shown in boldface type (for example, **tqbsp**).
- For Microsoft Windows systems, when a command or directory path contains spaces, the entire specification should be enclosed in double quotation (“ ”) unless noted otherwise. For example:

`"C:\Program Files\TeamQuest\manager"`

- In this manual, directory paths are shown using slash marks (/), as used in UNIX/Linux operating systems. For Microsoft Windows operating systems, you need to replace the slash marks (/) with backslashes (\).

Contents

About This Manual

Section 1. Introduction

1.1.	Hierarchical Structure for Statistics	1-2
1.1.1.	Parameter Hierarchy	1-3
1.1.2.	Table Field Hierarchy	1-5
1.1.3.	Alternate Statistic Classification	1-5
1.2.	Table Classification Scheme	1-6
1.3.	Types of Data Stored in Performance Databases	1-7
1.4.	How TeamQuest Manager Processes State Data	1-8

Section 2. AutoPredict Statistics

2.1.	Components of Response Statistics	2-2
2.2.	AutoPredict Growth Rate Statistics	2-6
2.3.	AutoPredict Indicator Statistics	2-8
2.4.	AutoPredict Status Statistics	2-11
2.5.	AutoPredict Resources Statistics	2-13

Section 3. DB2 Universal Database (UDB) Server

3.1.	DB2 Application Detail Statistics	3-3
3.2.	Bufferpool Statistics	3-17
3.3.	Database Configuration Statistics	3-23
3.4.	Database Statistics	3-24
3.5.	Database Status Statistics	3-41
3.6.	Instance Configuration Statistics	3-44
3.7.	Instance Statistics	3-45
3.8.	Instance Status Statistics	3-50
3.9.	Table Statistics	3-52
3.10.	Tablespace Statistics	3-54

Section 4. Hewlett-Packard HP-UX Systems

4.1.	System Activity Statistics	4-2
4.2.	Disk Space Statistics	4-21
4.3.	Network Statistics	4-25
4.4.	Workload Statistics	4-35
4.5.	Process Statistics	4-40
4.6.	Hardware Inventory Statistics	4-46
4.7.	System Log Statistics	4-52
4.8.	General Log Statistics	4-53
4.9.	TeamQuest Log Statistics	4-54
4.10.	Derived Statistics	4-55
4.10.1.	Workload Performance Derived Statistics	4-55
4.10.2.	TeamQuest On the Web Derived Statistics	4-57
4.10.3.	TeamQuest Alert Derived Statistics	4-59

Section 5. Hyper-V Statistics

5.1.	Hypervisor Statistics	5-2
5.2.	Virtual Machine Statistics	5-24
5.3.	Host Statistics	5-36
5.4.	I/O Statistics	5-53
5.5.	Network Statistics	5-56

Section 6. IBM AIX Systems

6.1.	System Activity Statistics	6-3
6.2.	Disk Space Statistics	6-34
6.3.	Network Statistics	6-38
6.4.	Workload Manager Statistics	6-51
6.5.	Workload Statistics	6-52
6.6.	LPAR Configuration Statistics	6-57
6.7.	Process Statistics	6-62
6.8.	Hardware Inventory Statistics	6-68
6.9.	System Log Statistics	6-74
6.10.	General Log Statistics	6-75
6.11.	TeamQuest Log Statistics	6-76
6.12.	Derived Statistics	6-77
6.12.1.	Workload Performance Derived Statistics	6-77
6.12.2.	TeamQuest On the Web Derived Statistics	6-79
6.12.3.	TeamQuest Alert Derived Statistics	6-80

Section 7. KVM Systems

7.1.	Libvirt.Host Config	7-3
7.2.	Libvirt.Host CPU Summary	7-4
7.3.	Libvirt.Host Device	7-5
7.4.	Libvirt.Host Memory Summary	7-6
7.5.	Libvirt.Host Network Config	7-7
7.6.	Libvirt.Host Network Filter	7-8
7.7.	Libvirt.Host Numa Memory	7-9
7.8.	Libvirt.Host Storage Pool Config	7-10
7.9.	Libvirt.Host Volume Config	7-11
7.10.	Libvirt.Virtual CPU	7-12
7.11.	Libvirt.Virtual Disk	7-14
7.12.	Libvirt.Virtual Disk Config	7-15
7.13.	Libvirt.Virtual Machine	7-16
7.14.	Libvirt.Virtual Machine Snapshot	7-18
7.15.	Libvirt.Virtual Memory	7-19
7.16.	Libvirt.Virtual Network Config	7-21
7.17.	Libvirt.Virtual Network Interface	7-22

Section 8. Linux Systems

8.1.	System Activity Statistics	8-3
8.2.	Disk Space Statistics	8-18
8.3.	Network Statistics	8-22
8.4.	LPAR Configuration Statistics	8-30
8.5.	Workload Statistics	8-32
8.6.	Process Statistics	8-35
8.7.	Hardware Inventory Statistics	8-43
8.8.	System Log Statistics	8-49
8.9.	General Log Statistics	8-50
8.10.	TeamQuest Log Statistics	8-51
8.11.	Derived Statistics	8-52
8.11.1.	Workload Performance Derived Statistics	8-52
8.11.2.	TeamQuest On the Web Derived Statistics	8-54
8.11.3.	TeamQuest Alert Derived Statistics	8-55

Section 9. Microsoft Windows Systems

9.1.	System Activity Statistics	9-2
9.2.	Workload Statistics	9-20
9.3.	Process Statistics	9-25
9.4.	Hardware Inventory Statistics	9-33
9.5.	General Log Statistics	9-40
9.6.	TeamQuest Log Statistics	9-41
9.7.	Windows Event Log Statistics	9-42
9.8.	Windows Services Statistics	9-43
9.9.	Derived Statistics	9-45
9.9.1.	TeamQuest Alert Derived Statistics	9-45
9.9.2.	TeamQuest On the Web Derived Statistics	9-46

9.9.3.	TeamQuest View Derived Statistics	9-47
9.10.	Optional System Activity Statistics	9-48
9.10.1.	Microsoft Exchange 2000	9-49
9.10.2.	Microsoft Exchange 5.5	9-52
9.10.3.	Microsoft IIS Statistics	9-54
9.10.4.	Microsoft SQL Server	9-57

Section 10. Network Applications

10.1.	Performance Statistics	10-2
10.2.	Client Statistics	10-7
10.3.	Port Statistics	10-9

Section 11. Network Devices

11.1.	Summary Statistics	11-2
11.2.	Interface Statistics	11-5

Section 12. Oracle Database Server

12.1.	Alert Log Alarm Statistics	12-3
12.2.	Block Contention Wait Statistics	12-4
12.3.	Datafile Capacity Statistics	12-5
12.4.	Datafile I/O Statistics	12-7
12.5.	Instance Alarm Statistics	12-9
12.6.	Instance Data Statistics	12-10
12.7.	Latch Statistics	12-37
12.8.	Library Cache Statistics	12-39
12.9.	Listener Alarm Statistics	12-41
12.10.	Lock Alarm Statistics	12-42
12.11.	Rollback Segment Statistics	12-43
12.12.	Row Cache Statistics	12-45
12.13.	Segment Alarm Statistics	12-47
12.14.	Session Statistics	12-49
12.15.	Session Wait Statistics	12-53
12.16.	System Parameters Statistics	12-55
12.17.	System Statistics	12-57
12.18.	System Wait Event Statistics	12-58
12.19.	Top SQL Cursors Statistics	12-60

Section 13. Oracle Solaris Systems

13.1.	System Activity Statistics	13-3
13.2.	Disk Space Statistics	13-28
13.3.	Network Statistics	13-32
13.4.	Workload Statistics	13-46
13.5.	Process Statistics	13-52
13.6.	Project Statistics	13-59

13.7.	Hardware Inventory Statistics	13-60
13.8.	System Log Statistics	13-68
13.9.	General Log Statistics	13-69
13.10.	TeamQuest Log Statistics	13-70
13.11.	Derived Statistics	13-71
13.11.1.	Workload Performance Derived Statistics	13-71
13.11.2.	TeamQuest On the Web Derived Statistics	13-73
13.11.3.	TeamQuest Alert Derived Statistics	13-75
13.11.4.	Rules Derived Statistics	13-76
13.12.	Zone Statistics	13-78
13.13.	Processor Set Statistics	13-80
13.14.	Resource Pool Statistics	13-81

Section 14. PostgreSQL Database Agent

14.1.	Database Summary Statistics	14-2
14.2.	Database Detail Statistics	14-3

Section 15. Sybase ASE Server

15.1.	Performance Statistics	15-3
15.2.	Active SQL Statistics	15-10
15.3.	Configuration Statistics	15-12
15.4.	Database Detail Statistics	15-14
15.5.	Database Summary Statistics	15-15
15.6.	Device Detail Statistics	15-16
15.7.	Engine Detail Statistics	15-17
15.8.	Lock Detail Statistics	15-19
15.9.	Process Statistics	15-21
15.10.	Process Waits Statistics	15-24
15.11.	System Waits Statistics	15-25

Section 16. System Alarm Statistics

Section 17. VMware Systems

17.1.	Physical Tables	17-2
17.1.1.	Block Device.by Host System Device Table	17-4
17.1.2.	Block Device.by Virtual Machine Table	17-7
17.1.3.	Block Device.VMware Summary Table	17-10
17.1.4.	CPU.by Host Processor Table	17-12
17.1.5.	CPU.by Virtual CPU Table	17-13
17.1.6.	CPU.by Virtual Machine Table	17-15
17.1.7.	CPU.by VMware Resource Table	17-17
17.1.8.	CPU.Relative Performance Table	17-18
17.1.9.	CPU.VMware Summary Table	17-19
17.1.10.	HINV.CPUModel Table	17-21
17.1.11.	HINV.CPU Thread Speeds Table	17-22

17.1.12.	HINV.Devices Table	17-23
17.1.13.	HINV.FileSystem Table	17-24
17.1.14.	HINV.Summary Table	17-25
17.1.15.	Memory.by Virtual Machine Table	17-27
17.1.16.	Memory.VMware Summary Table	17-30
17.1.17.	Network Device.by Host System Device Table ..	17-33
17.1.18.	Network Device.by Virtual Machine Table	17-35
17.1.19.	Network Device.vmnic by Virtual Machine Table	17-38
17.1.20.	Network Device.VMware Summary Table	17-39
17.1.21.	VMware.Availability by Virtual Machine Table	17-42
17.1.22.	VMware.Host Configuration Table	17-43
17.1.23.	VMware.Host Status Table	17-45
17.1.24.	VMware.Storage Configuration Table	17-47
17.1.25.	VMware.Support Metrics Table	17-48
17.1.26.	VMware.Virtual_Machines Table	17-49
17.1.27.	VMware Cluster.CPU Summary Table	17-51
17.1.28.	VMware Cluster.Memory Summary Table	17-53
17.1.29.	VMware Cluster.Resource Allocation Table	17-56
17.1.30.	VMware Cluster.Virtual Machine Operations	17-58
17.1.31.	VMware Datastore.File Type Usage by Datacenter Table	17-60
17.1.32.	VMware Datastore.Summary Table	17-62
17.1.33.	VMware Datastore.Usage by Virtual Machine Table	17-64
17.1.34.	VMware Resource Pool.CPU Summary Table ...	17-66
17.1.35.	VMware Resource Pool.Memory Summary Table	17-68
17.1.36.	VMware Resource Pool.Resource Allocation Table	17-71
17.1.37.	VMware Storage.Adapter by Host System Table	17-73
17.1.38.	VMware Storage.Adapter Summary Table	17-75
17.1.39.	VMware Storage.Datastore by Host System	17-76
17.1.40.	VMware Storage.Datastore by Virtual Machine ..	17-79
17.1.41.	VMware Storage.Datastore Summary	17-81
17.1.42.	VMware Storage.Path by Host System Table	17-82
17.1.43.	VMware Storage.Path Summary Table	17-84
17.1.44.	VMware Storage.Virtual Disk by Virtual Machine Table	17-85
17.2.	Derived Tables	17-87
17.2.1.	Cluster CPU Summary Table	17-89
17.2.2.	Cluster Memory Summary Table	17-91
17.2.3.	Cluster Resource Allocation Table	17-93
17.2.4.	Cluster Virtual Machine Operations Table	17-95
17.2.5.	Datastore Cluster File Type Usage Table	17-97
17.2.6.	Datastore Cluster Summary Table	17-99
17.2.7.	Datastore Cluster Usage by Virtual Machine Table	17-100
17.2.8.	Host Block Device Summary Table	17-102
17.2.9.	Host Block Device Usage Table	17-104
17.2.10.	Host CPU Resource Usage Table	17-107
17.2.11.	Host CPU Summary Table	17-108
17.2.12.	Host CPU Usage Table	17-110
17.2.13.	Host Memory Summary Table	17-111
17.2.14.	Host Network Device Summary Table	17-114

17.2.15.	Host Network Device Usage Table	17–116
17.2.16.	Resource Pool CPU Summary Table	17–118
17.2.17.	Resource Pool Memory Summary Table	17–120
17.2.18.	Resource Pool Resource Allocation Table	17–123
17.2.19.	Storage Adapter by Host System Table	17–125
17.2.20.	Storage Adapter Summary Table	17–127
17.2.21.	Storage Datastore by Host System Table	17–128
17.2.22.	Storage Datastore by Virtual Machine Table	17–131
17.2.23.	Storage Datastore Summary Table	17–133
17.2.24.	Storage Path by Host System Table	17–134
17.2.25.	Storage Path Summary Table	17–136
17.2.26.	Virtual Disk by Virtual Machine Table	17–137
17.2.27.	Virtual Machine Availability Table	17–139
17.2.28.	Virtual Machine Block Device Usage Table	17–141
17.2.29.	Virtual Machine Configuration Table	17–144
17.2.30.	Virtual Machine CPU Usage Table	17–146
17.2.31.	Virtual Machine Memory Usage Table	17–148
17.2.32.	Virtual Machine Network Device Usage Table	17–151
17.2.33.	Virtual Machine Virtual CPU Usage Table	17–153
17.2.34.	Virtual Machine Network vmnic Usage Table	17–155
17.2.35.	VMware CPU Relative Performance Table	17–157

Section 18. Web Server

18.1.	Web Server Statistics	18–2
18.2.	Top File Access Statistics	18–10

Bibliography	1
---------------------------	----------

Index	1
--------------------	----------

Section 1

Introduction

This section provides background information about the statistics stored in the TeamQuest performance database. It also provides a brief summary of the collection agents included with TeamQuest Manager.

This section contains the following topics:

- Hierarchical Structure for Statistics (see 1.1)
- Table Classification Scheme (see 1.2)
- Types of Data Stored in Performance Databases (see 1.3)
- How TeamQuest Manager Processes State Data (see 1.4)

1.1. Hierarchical Structure for Statistics

Data collected by the TeamQuest Manager agents is stored in the performance database in a hierarchical structure. The hierarchical levels consist of a statistic specification and its additional qualifiers. A statistic is defined by a hierarchy of key names that identify the statistic for retrieval from the database. The hierarchy of key names for a statistic is class, subclass, and statistic name.

For example, in the following statistic:

```
CPU:by Processor::%sys
```

CPU is the class; by Processor is the subclass; and %sys is the statistic name.

In general, statistics are grouped into two types related to the way each is stored. The statistic types differ in the kind of values that can be stored into each of them and in how they are qualified for a specific object. The two types of statistics are parameters and table fields.

- A parameter is a basic statistic stored in the aggregation sets of the database. Parameters can hold numeric values. A parameter has a fixed number of qualifying characteristics: system, resource (optional), and workload (optional). Parameters and aggregation sets are only supported with the TeamQuest database architecture. For information on the parameter hierarchy of key names, see 1.1.1.
- A table field is a statistic stored as a field in a table. A field can hold string or numeric values. Each table has a variable set of fields that identify the qualifying characteristics for each record. The remaining fields in the table contain performance metrics used to describe how the object is behaving. For information on the table field hierarchy of key names, see 1.1.2.

1.1.1. Parameter Hierarchy

A parameter is a basic statistic stored in the aggregation sets of the performance database. The term *parameter* refers to a fully qualified statistic, which consists of the statistic and its additional qualifiers (a system specification, an optional resource specification, and an optional workload specification).

The hierarchy of key names defines a parameter and has the following format with each level of the hierarchy separated by a colon (:).

```
System:Class:Subclass:Statistic:Resource:Workload Set:Workload
```

Statistic Specification

A statistic describes the information being collected and stored in the database.

The following shows the hierarchy of key names for several statistics collected by the TeamQuest Manager agents:

Class	Subclass	Statistic Name
Block Device	by Device:	avserv
CPU	by Processor:	%sys
Kernel	Buffers:	bread/s
Memory		freemem
RPC	Client:Connectionless	timeouts/s

The hierarchy of key names that is used for the classification of a statistic are the following:

- **Class** identifies a set of performance data with a general relationship. Any given system can have many classes and the class key name is used to group related subclasses. Examples of a class include Block Device, CPU, Kernel, Memory, and RPC.
- **Subclass** divides a class into a collection of related data. For example, the class Block Device contains subclasses called by Device and Summary.

A subclass can consist of two levels. Either one or both levels can be used in the subclass name, or both levels can be empty. For example, the by Processor: subclass of the CPU class has a single level. The Client:Connectionless subclass of the RPC class has two levels. The Memory class has no subclass associated with it and both levels are empty. In the subclass key name, the two levels are separated by a colon (:). When a subclass has only one level, the second level of the subclass key name in the parameter hierarchy is empty.

A class can also have no subclasses associated with it. When this occurs, the parameter hierarchy contains :: for the subclass key position.

Introduction

- **Statistic name** is the name of the information being collected and stored in the database. For example, avserv, %sys, bread/s, freemem, and timeouts/s are statistics names.

Using the parameter format, the statistics in the previous table (assuming they were collected on a system called saturn) would appear as follows:

```
saturn:Block Device:by Device::avserv::  
saturn:CPU:by Processor::%sys::  
saturn:Kernel:Buffers::bread/s::  
saturn:Memory::freemem::  
saturn:RPC:Client:Connectionless:timeouts/s::
```

Parameter Qualifiers

The following qualifiers are used for parameters stored in the performance database:

- The **system** qualifier specifies the system on which the data was collected. A parameter must be qualified by a system name.
- A **resource** is an optional qualifier for a parameter. A resource is a specific object relating to the class and subclass to which it belongs. A resource could be a physical resource, such as a disk or tape unit, or it could be a logical resource, such as a system call. For example, a resource for disk parameters would be the name of a particular disk. In this case, a separate parameter for each disk on the system is stored in the database.

A statistic with a hierarchy of Block Device:by Device:avserv can be further qualified by the specific disk unit of disk-0. Using the parameter format, the parameter hierarchy for this statistic would appear as:

```
saturn:Block Device:by Device::avserv:disk-0::
```

When a resource is not being used as a qualifier, the parameter hierarchy contains :: indicating that its key position is empty.

- A **workload** is an optional qualifier for a parameter. A workload is a logical classification of the work performed on the system. The workload classifications are site-specific and are expressed in terms that are meaningful to the business entity. For example, by defining the appropriate workloads, the CPU resources on the system could be summarized by department name (payroll, development, or marketing).

A workload qualifier consists of a workload set name and a workload. Using the parameter format, the parameter hierarchy for a statistic with a workload qualifier would appear as follows. The workload set is Example and the workload is users.

```
saturn:Workload:by Workload::%cpu::Example:users
```

When a workload is not being used as a qualifier, the parameter hierarchy contains :: indicating that its key position is empty.

1.1.2. Table Field Hierarchy

A table field is a statistic stored as a field in a table. A table is a collection of records containing columns of data that are related in some way. Each record in a table contains data collected at a given time to capture a history of performance information. A record contains one set of fields (for example, column values) for the table.

The table name is specified by the class and subclass key names (for example, System.Alarms and HINV.Summary).

A table field is defined by a hierarchy of key names that identify the data for retrieval from the database. The key names used for a table field are class, subclass, and a statistic name (for example, System.Alarms.severity and HINV.Summary.CPU Type).

- **Class** identifies a general grouping of table information. Examples of a class include HINV, Oracle, Web, and System.
- **Subclass** identifies a group of table fields within the table class. Examples of subclasses within the Oracle class are LibraryCache, Rollback, and SystemWait.
- **Statistic name** identifies the information being collected and stored in the table. Examples of statistics stored in the Oracle.LibraryCache table are Actual_Interval, get_hits, gets, and reloads.

Within each table stored in the TeamQuest performance database, certain fields are designated as identifiers. Identifiers are selected fields that are used as the qualifying characteristics for records. For example, the combination of Timestamp, Sequence_Number, and System is the unique identifier for an alarm in the System.Alarms table.

Each table record contains a qualifier called System which records the system name associated with the data collected.

1.1.3. Alternate Statistic Classification

The Class, Subclass, and Statistic Name classification of statistics is used in most of the TeamQuest Performance Software (TeamQuest View, and TeamQuest Manager). However, some utilities may still use the previous classification scheme of earlier versions of the software. This previous classification scheme consisted of Category Group, Category, Subcategory, and Statistic Name for parameters; and Table Class, Table Name, and Field Name for table fields.

The mapping of the older classification to the new classification is as follows:

Category Group -> Class
Category.Subcategory -> Subclass
Statistic Name -> Statistic Name

Table Class -> Class
Table Name -> Subclass
Field Name -> Statistic Name

1.2. Table Classification Scheme

TeamQuest Analyzer and the open database architecture use a different classification scheme for referencing statistics. Each statistic, which is known as a field, is grouped into an associated table. A table name may contain one or more logical levels, each delimited by a period (.). The following are examples of table names:

SOLARIS.Process
Oracle.Session
Processor
DB2.Database.Bufferpool Summary

If any part of the table name contains a literal period (.) or a backslash (\), the character must be preceded by a backslash (\). For example, the following table names contain a period in the name and, therefore, must have a backslash before the period:

\.Net CLR Memory
ASP\NET

Statistics in this manual, which are documented using the Class:Subclass notation, are mapped into the table/field classification scheme by converting the colons (:) to periods (.) wherever they appear.

Parameters

Table Name: Processor
Fields: Time, System, Available Bytes, Committed Bytes, and so on

Table Name: DB2.Database.Bufferpool Summary
Fields: Time, System, Instance, Resource, LogicalDataReads/s, PhysicalDataReads/s, and so on

Table Fields

Table Name: SOLARIS.Process
Fields: Time, System, command, totcpu, and so on


Table Name: Oracle.Session
Fields: Time, System, Instance, sid, ora_user, cpu_used, sql_text, and so on

1.3. Types of Data Stored in Performance Databases

The TeamQuest Manager agents collect performance, state, and event data. This data is stored in performance databases.

Performance Data

Performance (intervalized) data is sampled and recorded at regular time intervals. An example of performance data is a record from the CPU.by Processor table. Records from this table are typically displayed in the form of a graph with one or more data series displayed across time. The example area graph below displays four different CPU utilization statistics over a day.



State Data

State data represents a state that existed over multiple time intervals. State data is checked at regular intervals and is recorded only when the state has changed. State data may also be recorded when the collection agents are started or restarted. State data is never aggregated or consolidated. An example of state data is a record from the HINV.Summary table. Records from this table are typically displayed in the form of a table where each record represents the time and date when state information changed. The example table below shows the state of the hardware inventory information for the time period of the three systems.

New Chart 08/18/14 00:00 - 08/18/14 16:00 3/3				
Time	System	Logical CPU Count	Memory	OS Release
08/18/14 00:00	SKROEMER	2	3 GBytes	Windows XP 5.01
08/18/14 19:01	clsuse1x64	2	1.97 GBytes	Linux 2.6.16.60-0
08/18/14 00:00	CTHORPE	2	4 GBytes	Windows XP 5.01

Introduction

Event Data

Event data is a record of the occurrence of events. Event records are created when events are detected. Event detection occurs at regular intervals. The time that an event occurred is saved along with the time the event was detected. Event data is never aggregated or consolidated. An example of event data is a record from the System.Alarms table. Records from this table are typically displayed in the form of a table. The example table below shows all the alarms that were triggered on a server in a 24-hour period.

Time	System	Severity	Alarm ID
08/18/14 16:09	CLW2K3-6	Critical	High_Paging_Rate_Alarm
08/18/14 16:24	CLW2K3-6	Critical	Busy_Disk_Alarm [Total]
08/18/14 16:33	CLW2K3-6	Normal	High_Paging_Rate_Alarm
08/18/14 16:43	CLW2K3-6	Normal	Busy_Disk_Alarm [Total]
08/18/14 16:52	CLW2K3-6	Normal	Busy_Disk_Alarm [1 E:]
08/18/14 17:22	CLW2K3-6	Critical	Busy_Disk_Alarm [1 E:]
08/18/14 17:56	CLW2K3-6	Normal	Busy_Disk_Alarm [1 E:]
08/18/14 18:26	CLW2K3-6	Warning	Busy_Disk_Alarm [1 E:]
08/18/14 19:56	CLW2K3-6	Normal	Busy_Disk_Alarm [1 E:]
08/18/14 20:09	CLW2K3-6	Critical	High_Paging_Rate_Alarm
08/18/14 20:13	CLW2K3-6	Normal	High_Paging_Rate_Alarm
08/18/14 20:26	CLW2K3-6	Warning	Busy_Disk_Alarm [1 E:]
08/18/14 21:16	CLW2K3-6	Normal	Busy_Disk_Alarm [1 E:]

1.4. How TeamQuest Manager Processes State Data

When selecting performance data or event data from a TeamQuest database, TeamQuest Manager retrieves data records that were recorded with timestamps within a requested time interval. When state data records are recorded, the state being recorded remains in effect until new state data records are recorded. Even though state data may have been recorded before a specified time interval, the previously recorded state remains in effect at the start of the specified interval.

When retrieving data for TeamQuest Analyzer, **tqgetm**, TeamQuest tView, or **tqharvc_open**, the state data recorded most recently before the start of the selected interval are returned. In this way, records reflecting the state during the entire selected interval are returned, even if some of the state data were recorded before the start of the interval.

For database performance reasons, limitations have been established on how far back in time TeamQuest Manager searches for previously recorded state data records. When searching backward for the state data recorded most recently before the start of the selected time interval, TeamQuest Manager searches up to 73 hours before the start of the selected time interval. If no appropriate state data is found in the 73-hour period, no additional records are returned. This limitation is not enforced when retrieving state data from an open database architecture.

As an example, HINV.Summary records are typically recorded every day at midnight. In TeamQuest Analyzer, if you want to report on the hardware inventory and select a time period of 8:00 to 12:00, you will see the HINV.Summary data recorded at midnight, even though it was not recorded during the time period of 8:00 to 12:00 because the data retrieval handles state data differently from performance data or event data. If no HINV.Summary data has been recorded for the past 4 days, no records will be returned because the record retrieval does not search backward more than 73 hours before the requested start time.

The following TeamQuest Manager tables are examples of state tables to which this processing applies:

- HINV.Devices
- HINV.FileSystem
- HINV.Summary
- VMware.Host Configuration

Section 2

AutoPredict Statistics

The AutoPredict Agent is a model-based analysis tool. Based on user-specified policies, the AutoPredict Agent automatically builds, calibrates and solves models, and stores the results into a TeamQuest performance database. The AutoPredict Agent provides information to perform TeamQuest Performance Indicator, stretch factor, and components of response analysis on your system.

This section contains a listing of the statistics collected by the agent:

- Components of Response Statistics (see 2.1)
- AutoPredict Growth Rate Statistics (see 2.2)
- AutoPredict Indicator Statistics (see 2.3)
- AutoPredict Status Statistics (see 2.4)
- AutoPredict Resources Statistics (see 2.5)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

2.1. Components of Response Statistics

The following statistics are stored in the TeamQuest performance database by the AutoPredict Agent. These statistics are used to detect where the workload spends time on the server.

Components of response statistics represent the time CPU and I/O devices are used by the workload. These components are estimates of the time a workload is active at the resource (service time) and the time a workload is waiting for the resource (queue delay).

Table Field Hierarchy

Class:	Components of Response
Subclass:	by Workload
IT Resource Name:	TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	Components of Response.by Workload
Open Table Name:	COMPOFRSPBYWL
Collection interval:	Based on the collection period
Default retention:	1 month
Table type:	Performance

Statistic Name	Description
Actual_Interval	The actual period of time the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
AutoPredict_Interval	The name of the interval used by the AutoPredict Agent [Sequential = ID Non-Sequential = ID]
Days_to_Step	The point in the future represented by this prediction step. This is the number of days from the timestamp of the data sample being analyzed. [Sequential = ID Non-Sequential = ID]
Frame	The name of the frame for the model. When a growth scenario is used, this is Growth. If no growth scenario is used, this is Frame 1. [Sequential = ID Non-Sequential = ID]
Interval	The expected duration of the interval that the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
Logical_System	The name of the logical system or partition. This field is limited to 80 characters. Any Logical_System name longer than 80 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Physical_System	The name of the physical system or host system. This field is limited to 52 characters. Any Physical_System name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Queue_Delay	The amount of time the workload was waiting for the resource [Sequential = AVG Non-Sequential = NON]

Resource	The name of the physical device. This field is limited to 51 characters. Any Resource name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Resource_Type	The type of resource or physical device, such as CPU or disk unit. This field is limited to 51 characters. Any Resource_Type longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Sequence	The relative order of the various components of response. The CPU service time and CPU delay time are ordered first, followed by the I/O devices in order of their contribution to response time. [Sequential = AVG Non-Sequential = AVG]
Service	The amount of time the workload was active or in service at the resource [Sequential = AVG Non-Sequential = NON]
Step	The name of the step that was solved. When a growth scenario is used, this is a string that identifies the period being predicted. The period is the number of days into the future. A value of 0 indicates the baseline step. If no growth scenario is used, this is Step 1. [Sequential = ID Non-Sequential = ID]
Step_Date	The timestamp of the point in the future represented by this prediction step [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Workload	The name of the workload. A workload represents a logical classification of work performed on the system. [Sequential = ID Non-Sequential = ID]
Workload_Set	The name of the workload set [Sequential = ID Non-Sequential = ID]

Table Field Hierarchy

Class:	Components of Response
Subclass:	Summary
IT Resource Name:	TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	Components of Response.Summary
Open Table Name:	COMPOFRSPSUM
Collection interval:	Based on the collection period
Default retention:	1 month
Table type:	Performance

Statistic Name	Description
Actual_Interval	The actual period of time the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
AutoPredict_Interval	The name of the interval used by the AutoPredict Agent [Sequential = ID Non-Sequential = ID]
Days_to_Step	The point in the future represented by this prediction step. This is the number of days from the timestamp of the data sample being analyzed. [Sequential = ID Non-Sequential = ID]
Frame	The name of the frame for the model. When a growth scenario is used, this is Growth. If no growth scenario is used, this is Frame 1. [Sequential = ID Non-Sequential = ID]
Interval	The expected duration of the interval that the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
Logical_System	The name of the logical system or partition. This field is limited to 80 characters. Any Logical_System name longer than 80 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Physical_System	The name of the physical system or host system. This field is limited to 52 characters. Any Physical_System name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Queue_Delay	The amount of time the workload was waiting for the resource [Sequential = AVG Non-Sequential = NON]
Resource	The name of the physical device. This field is limited to 51 characters. Any Resource name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Resource_Type	The type of resource or physical device, such as CPU or disk unit. This field is limited to 51 characters. Any Resource_Type longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]

Sequence	The relative order of the various components of response. The CPU service time and CPU delay time are ordered first, followed by the I/O devices in order of their contribution to response time. [Sequential = AVG Non-Sequential = AVG]
Service	The amount of time the workload was active or in service at the resource [Sequential = AVG Non-Sequential = NON]
Step	The name of the step that was solved. When a growth scenario is used, this is a string that identifies the period being predicted. The period is the number of days into the future. A value of 0 indicates the baseline step. If no growth scenario is used, this is Step 1. [Sequential = ID Non-Sequential = ID]
Step_Date	The timestamp of the point in the future represented by this prediction step [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Workload_Set	The name of the workload set. [Sequential = ID Non-Sequential = ID]

2.2. AutoPredict Growth Rate Statistics

The following statistics are stored in the TeamQuest performance database by the AutoPredict Agent. One record is written when a new growth rate is calculated. If the growth rate analysis period is a day, a record is written each day. If the growth rate analysis period is a week, a record is written each week. If the growth rate analysis period is a month, a record is written each month.

Table Field Hierarchy

Class:	AutoPredict
Subclass:	Growth Rate
IT Resource Name:	TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	AutoPredict.Growth Rate
Open Table Name:	APGROWTH
Collection interval:	Based on the collection period
Default retention:	5 years
Table type:	Performance

Statistic Name	Description
AutoPredict_Growth_Rate	The name of the AutoPredict Growth Rate object as named in the TeamQuest Administration Console. This field is limited to 80 characters. [Sequential = ID Non-Sequential = ID]
Evaluation Statistic	The statistic name that is used to calculate the growth rate. The format of this field is <table name>:<statistic name>. [Sequential = ID Non-Sequential = ID]
Goodness_of_Fit	The goodness of fit describes how well the calculated growth rate fits the set of observations. This field can be any number between 0.0 and 1.0, where 0.0 is the worst goodness of fit and 1.0 is the best goodness of fit. [Sequential = ID Non-Sequential = ID]
Growth_Rate	The annualized growth rate. This growth rate is used for calculating the growth per step. This field can be any number between -100.00 percent and a very large percentage. [Sequential = ID Non-Sequential = ID]
Growth_Rate_Type	The growth rate type is always a compound percentage. [Sequential = ID Non-Sequential = ID]
Logical_System	The name of the logical system or partition. This field is limited to 80 characters. Any Logical_System name longer than 80 characters is truncated. [Sequential = ID Non-Sequential = ID]
Physical_System	The name of the physical system or host system. This field is limited to 52 characters. Any Physical_System name longer than 52 characters is truncated. [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Where_Clause	A where clause that is supplied by the user in the TeamQuest Administration Console. If no where clause is supplied, this field is N/A. For more information on where clauses, see the section on working with policies in the <i>TeamQuest Administration Console User Guide</i> . [Sequential = ID Non-Sequential = ID]
Workload	The name of the workload. A workload represents a logical classification of work performed on the system. If no workload is supplied, this field is N/A. [Sequential = ID Non-Sequential = ID]
Workload_Set	The name of the workload set. If no workload is supplied, this field is N/A. [Sequential = ID Non-Sequential = ID]

2.3. AutoPredict Indicator Statistics

The following statistics are stored in the TeamQuest performance database by the AutoPredict Agent.

The TeamQuest Performance Indicator (TPI) is an expression of the comparison between the actual performance of a workload and the theoretical optimal performance. The TPI is an important value for indicating bottlenecks in your system. It is calculated by dividing the total service time by the total response time (service time + queue delay) times 100. If the queue delay at a resource increases, the value for TPI decreases. In a perfect system with no queuing, the value for TPI is 100. In a system where the service time and queue delay are equal, a value for TPI is 50. A value less than 50 indicates you may be experiencing more queue delay than service and indicates a possible bottleneck.

Table Field Hierarchy

Class:	AutoPredict Indicators
Subclass:	by Workload
IT Resource Name:	TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	AutoPredict Indicators.by Workload
Open Table Name:	APINDBYWL
Collection interval:	Based on the collection period
Default retention:	1 month
Table type:	Performance

Statistic Name	Description
Actual_Interval	The actual period of time the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
AutoPredict_Interval	The name of the interval used by the AutoPredict Agent [Sequential = ID Non-Sequential = ID]
Capacity_Rule_Exceeded	When one or more maximum capacity rules are exceeded, this field contains a description of a rule that was exceeded. If no maximum capacity rules are exceeded, this field contains <N/A>. The format of the maximum capacity rule description is <rule name> <rule component> < predicted value comparison>. [Sequential = ID Non-Sequential = ID]
Days_to_Step	The point in the future represented by this prediction step. This is the number of days from the timestamp of the data sample being analyzed. [Sequential = ID Non-Sequential = ID]
Frame	The name of the frame for the model. When a growth scenario is used, this is Growth. If no growth scenario is used, this is Frame 1. [Sequential = ID Non-Sequential = ID]
Interval	The expected duration of the interval that the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]

Logical_System	The name of the logical system or partition. This field is limited to 80 characters. Any Logical_System name longer than 80 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Physical_System	The name of the physical system or host system. This field is limited to 52 characters. Any Physical_System name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Response_Time	The predicted average response time for a transaction for the workload, in seconds [Sequential = AVG Non-Sequential = AVG]
Step	The name of the step that was solved. When a growth scenario is used, this is a string that identifies the period being predicted. The period is the number of days into the future. A value of 0 indicates the baseline step. If no growth scenario is used, this is Step 1. [Sequential = ID Non-Sequential = ID]
Step_Date	The timestamp of the point in the future represented by this prediction step [Sequential = ID Non-Sequential = ID]
Stretch_Factor	The stretch factor for the workload [Sequential = DIV Non-Sequential = DIV]
Throughput	The predicted throughput for the workload, in units of transactions per second Sequential = AVG Non-Sequential = AVG]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
TPI	The TeamQuest Performance Indicator (TPI) for the workload. [Sequential = DIV Non-Sequential = DIV]
Workload	The name of the workload. A workload represents a logical classification of work performed on the system. [Sequential = ID Non-Sequential = ID]
Workload_Set	The name of the workload set [Sequential = ID Non-Sequential = ID]

AutoPredict Statistics

Table Field Hierarchy

Class:	AutoPredict Indicators
Subclass:	Summary
IT Resource Name:	TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	AutoPredict Indicators.Summary
Open Table Name:	APINDSUM
Collection interval:	Based on the collection period
Default retention:	1 month
Table type:	Performance

Statistic Name	Description
Actual_Interval	The actual period of time the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
AutoPredict_Interval	The name of the interval used by the AutoPredict Agent [Sequential = ID Non-Sequential = ID]
Days_to_Step	The point in the future represented by this prediction step. This is the number of days from the timestamp of the data sample being analyzed. [Sequential = ID Non-Sequential = ID]
Frame	The name of the frame for the model. When a growth scenario is used, this is Growth. If no growth scenario is used, this is Frame 1. [Sequential = ID Non-Sequential = ID]
Interval	The expected duration of the interval that the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
Logical_System	The name of the logical system or partition. This field is limited to 80 characters. Any Logical_System name longer than 80 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Physical_System	The name of the physical system or host system. This field is limited to 80 characters. Any Physical_System name longer than 80 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Step	The name of the step that was solved. When a growth scenario is used, this is a string that identifies the period being predicted. The period is the number of days into the future. A value of 0 indicates the baseline step. If no growth scenario is used, this is Step 1. [Sequential = ID Non-Sequential = ID]
Step_Date	The timestamp of the point in the future represented by this prediction step [Sequential = ID Non-Sequential = ID]
Stretch_Factor	The stretch factor for the work on the system [Sequential = DIV Non-Sequential = DIV]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
TPI	The TeamQuest Performance Indicator (TPI) for the workload. [Sequential = DIV Non-Sequential = DIV]
Workload_Set	The name of the workload set [Sequential = ID Non-Sequential = ID]

2.4. AutoPredict Status Statistics

The following statistics are stored in the TeamQuest performance database by the AutoPredict Agent. One record is written per model, and it contains the summary for all steps. These statistics are the results of the automatic evaluation of model results.

Table Field Hierarchy

Class:	AutoPredict
Subclass:	Status
IT Resource Name:	TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	AutoPredict.Status
Open Table Name:	APSTATUS
Collection interval:	Based on the collection period
Default retention:	1 month
Table type:	Performance

Statistic Name	Description
Actual_Interval	The actual period of time represented by the model built and stored by the AutoPredict Agent [Sequential = SUM Non-Sequential = ID]
AutoPredict_Interval	The name of the interval used by the AutoPredict Agent [Sequential = ID Non-Sequential = ID]
Capacity_Rule_Exceeded	When one or more maximum capacity rules are exceeded, this field contains a description of a rule that was exceeded. If no maximum capacity rules are exceeded or depending on the Solver_Status value, this field contains <N/A>. The format of the maximum capacity rule description is <rule name> <rule component> <predicted value comparison>. [Sequential = ID Non-Sequential = ID]
Compliance_Status	The status of compliance with capacity rules. "Noncompliance" indicates at least one capacity rule is out of compliance. "Compliance" indicates all capacity rules are in compliance. "No rule" indicates that there were no compliance rules to evaluate. This field may contain <N/A>, depending on the Solver_Status value. [Sequential = ID Non-Sequential = ID]

AutoPredict Statistics

Critical_Resource	<p>When one or more maximum capacity rules are exceeded, this field contains the name of the resource where the workload spends most of its time. If no maximum capacity rules are exceeded or depending on the Solver_Status value, this field contains <N/A>. This will take its data from the Components of Response Summary table.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Days_To_Noncompliance	<p>The number of days from the baseline model date when the system goes to a noncompliance state. A value of -1 indicates that the system is in compliance for all modeled periods. This field may contain <N/A>, depending on the Solver_Status value.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Frame	<p>The name of the frame for the model. When a growth scenario is used, this is Growth. If no growth scenario is used, this is Frame 1.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Interval	<p>The expected duration of the interval that the AutoPredict Agent builds and solves models for</p> <p>[Sequential = SUM Non-Sequential = ID]</p>
Logical_System	<p>The name of the logical system or partition. This field is limited to 80 characters. Any Logical_System name longer than 80 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Noncompliance_Date	<p>The predicted date when the system goes to a noncompliance state. A value of zero indicates that the system is in compliance for all modeled periods. This field may contain <N/A>, depending on the Solver_Status value.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Physical_System	<p>The name of the physical system or host system. This field is limited to 52 characters. Any Physical_System name longer than 52 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Solver_Status	<p>The status of the Solver. The status can be any of the following:</p> <ul style="list-style-type: none">Success = The model solved successfully.Data Not Available = The data needed to solve the model could not be found.Calibration Failure = The model could not be calibrated.Calibration Timeout = The calibration took too long to complete.Solve Step <Step Name> Failed = The step with the given name failed.Solve Step <Step Name> Timeout = The step with the given name took too long to solve. <p>If the Solver_Status value is Data Not Available, Calibration Failure, Calibration Timeout, or if the very first step fails or times out, the values for compliance related fields will be <N/A>.</p> <p>If the Solver_Status value is Success or if a step greater than 1 has failed or timed out, the compliance related fields will be populated with data from the successful steps.</p> <p>[Sequential = ID Non-Sequential = ID]</p>

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Workload_Set	The name of the workload set [Sequential = ID Non-Sequential = ID]

2.5. AutoPredict Resources Statistics

The following statistics are stored in the TeamQuest performance database by the AutoPredict Agent. Some of these statistics are for automatic evaluation of modeling results. These statistics will most likely be used for manual, after the fact, analysis.

Table Field Hierarchy

Class:	AutoPredict
Subclass:	Resources
IT Resource Name:	TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	AutoPredict.Resources
Open Table Name:	APRESOUR
Collection interval:	Based on the collection period
Default retention:	1 month
Table type:	Performance

Statistic Name	Description
Actual_Interval	The actual period of time the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
AutoPredict_Interval	The name of the interval used by the AutoPredict Agent [Sequential = ID Non-Sequential = ID]
Days_to_Step	The point in the future represented by this prediction step. This is the number of days from the timestamp of the data sample being analyzed. [Sequential = ID Non-Sequential = ID]
Frame	The name of the frame for the model. When a growth scenario is used, this is Growth. If no growth scenario is used, this is Frame 1. [Sequential = ID Non-Sequential = ID]
Interval	The expected duration of the interval that the AutoPredict Agent builds and solves models for [Sequential = SUM Non-Sequential = ID]
Logical_System	The name of the logical system or partition. This field is limited to 80 characters. Any Logical_System name longer than 80 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Percent_Busy	The predicted percent utilization of the resource [Sequential = AVG Non-Sequential = AVG]

AutoPredict Statistics

Physical_System	The name of the physical system or host system. This field is limited to 52 characters. Any Physical_System name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Queue_Length	The predicted queue length (in service plus waiting) of the resource [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the physical device. This field is limited to 51 characters. Any Resource name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Resource_Type	The type of resource or physical device, such as CPU or disk unit. This field is limited to 51 characters. Any Resource_Type longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Step	The name of the step that was solved. When a growth scenario is used, this is a string that identifies the period being predicted. The period is the number of days into the future. A value of 0 indicates the baseline step. If no growth scenario is used, this is Step 1. [Sequential = ID Non-Sequential = ID]
Step_Date	The timestamp of the point in the future represented by this prediction step [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Workload_Set	The name of the workload set [Sequential = ID Non-Sequential = ID]

Section 3

DB2 Universal Database (UDB) Server

The DB2 UDB Agent collects information on DB2 Universal Database (UDB) instances. The agent obtains instance summary data pertaining to sorts, connections, and agents. It also gathers database summary data relating to connections, agents, lock summary information, bufferpool I/O, sorts, SQL counts, row I/O counts, log space usage, and memory usage. Additional detailed information about bufferpools, tablespaces, tables, applications, instance status, instance configuration, database status, and database configuration is retrieved by the agent as well.

This section contains a listing of the statistics collected by the agent:

- DB2 Application Detail Statistics (see 3.1)
- Bufferpool Statistics (see 3.2)
- Database Configuration Statistics (see 3.3)
- Database Statistics (see 3.4)
- Database Status Statistics (see 3.5)
- Instance Configuration Statistics (see 3.6)
- Instance Statistics (see 3.7)
- Instance Status Statistics (see 3.8)
- Table Statistics (see 3.9)
- Tablespace Statistics (see 3.10)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

In addition, the end of each statistic description contains a notation in brackets relating to any DB2 monitor switches that need to be enabled in order to collect the particular statistic (for example, [Switch = table]). This notation is in the form [Switch = switchname], where switchname is the name of the DB2 switch needed. The following switch names are used: none, bufferpool, lock, sort, statement, table, and uow.

3.1. DB2 Application Detail Statistics

Detailed information about the applications for each DB2 instance is stored within the DB2.Application Detail table in the TeamQuest performance database.

Table Field Hierarchy

Class:	DB2
Subclass:	Application Detail
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Application Detail
Open Table Name:	DB2APPDETAIL
Collection interval:	1 minute (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
acc_cur_reqs	The number of times that a request by the application for an I/O block was accepted during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the application started within the given sample interval. [Sequential = SUM Non-Sequential = ID] [Switch = none]
agent_id	The agent identifier of the agent serving the application [Sequential = ID Non-Sequential = ID] [Switch = none]
agent_pid	The process identifier (UNIX/Linux systems) or thread identifier (Windows systems) of a DB2 UDB Agent [Sequential = ID Non-Sequential = ID] [Switch = none]
agent_sys_cpu	The amount of time in seconds spent by the agents for this application executing system calls during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
agent_usr_cpu	The amount of time in seconds spent by the agents for this application executing database manager code during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
agents_stolen	The number of times an agent was stolen from the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]

appl_id	The unique identifier for the application [Sequential = ID Non-Sequential = ID] [Switch = none]
appl_name	The name of the application running on the client as known to the database manager [Sequential = ID Non-Sequential = ID] [Switch = none]
appl_status	The current status of the application [Sequential = LST Non-Sequential = ID] [Switch = none]
assoc_agents	The number of subagents associated with the application [Sequential = LST Non-Sequential = SUM] [Switch = none]
auth_id	The authorization identifier of the user who invoked the application that is being monitored [Sequential = ID Non-Sequential = ID] [Switch = none]
avg_d_rds	The average number of sectors that are read per direct read request during the interval. Not stored by default. [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_rdtype	The average time in milliseconds spent direct reading a sector by the application during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_wrts	The average number of sectors that are written per direct write request during the interval. Not stored by default. [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_wrttime	The average time in milliseconds spent direct writing a sector by the application during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_lckwtime	The average amount of time in milliseconds spent by the application waiting for a lock during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = lock]
avg_rdtype	The average time in milliseconds spent reading a page by the application during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_sort_time	The average amount of time in milliseconds per sort by the application during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = sort]

avg_wrttime	The average time in milliseconds spent writing a page by the application during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
binds	The number of binds and pre-compiles attempted by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
catc_heap_full	The number of times an insert into the catalog cache for the application failed due to a heap-full condition in the database heap during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
catc_inserts	The number of times the system tried to insert table descriptor information into the catalog cache for the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
catc_lookups	The number of times the catalog cache was referenced by the application to obtain table descriptor information during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
catc_max_mem	The maximum amount of memory in kilobytes used in the catalog cache by the application [Sequential = MAX Non-Sequential = SUM] [Switch = none]
catc_overflows	The number of times an insert into the catalog cache for the application failed due to the catalog cache being full during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
clnt_addr	The communication address of the client [Sequential = ID Non-Sequential = ID] [Switch = none]
clnt_comm	The communication protocol that the client application is using to communicate with the server. Not stored by default. [Sequential = ID Non-Sequential = ID] [Switch = none]
clnt_nname	The node name in the database manager configuration file at the client node [Sequential = ID Non-Sequential = ID] [Switch = none]
clnt_pid	The process identifier of the client application that made the connection to the database [Sequential = ID Non-Sequential = ID] [Switch = none]

clnt_platform	The operating system on which the client application is running [Sequential = ID Non-Sequential = ID] [Switch = none]
clnt_prdid	The product and version that is running on the client [Sequential = ID Non-Sequential = ID] [Switch = none]
clnt_userid	The user id that the user specified when logging into the operation system [Sequential = ID Non-Sequential = ID] [Switch = none]
codepage_id	The code page identifier for the application [Sequential = ID Non-Sequential = ID] [Switch = none]
commit_sqls	The number of commit SQL statements attempted by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
conn_end_time	The timestamp that the application completed a connection request [Sequential = LST Non-Sequential = ID] [Switch = none]
conn_start_time	The timestamp that the application started a connection request [Sequential = FST Non-Sequential = ID] [Switch = none]
coord_agent_pid	The process identifier (UNIX/Linux systems) or thread identifier (Windows systems) of the coordinator agent for the application [Sequential = ID Non-Sequential = ID] [Switch = none]
country_code	The country code of the database for which the monitor data is collected [Sequential = ID Non-Sequential = ID] [Switch = none]
d_rd_reqs	The number of requests to perform direct read operations by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_rds	The number of read operations by the application that do not use a bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_rdtype	The amount of time in milliseconds spent direct reading a sector by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_wrt_reqs	The number of requests to perform direct write operations by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]

d_wrts	The number of write operations by the application that do not use a bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_wrttime	The amount of time in milliseconds spent direct writing a sector by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
data_wrts	The number of data pages written by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
database	The name of the database to which the application is connected [Sequential = ID Non-Sequential = ID] [Switch = none]
ddl_sqls	The number of data definition language (DDL) SQL statements executed by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
deadlocks	The number of deadlocks by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = lock]
dynamic_sqls	The number of dynamic SQL statements attempted by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
eff_lock_waits	The number of times the application was waiting for a lock at the beginning of the sample interval or started waiting for a lock during the sample interval. [Sequential = SUM Non-Sequential = SUM] [Switch = lock]
exec_time	The current amount of time in milliseconds spent executing statements from the application [Sequential = LST Non-Sequential = SUM] [Switch = statement]
failed_sqls	The number of failed SQL statements by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
hit_ratio	The percentage of logical reads that were satisfied for the application without having to invoke a read from disk during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]

idle_time	The amount of time in seconds since the application last issued a request to the server [Sequential = LST Non-Sequential = ID] [Switch = statement]
idx_wrts	The number of index pages written by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
Instance	The name of the instance from which data is obtained. Up to 24 characters are displayed. [Sequential = ID Non-Sequential = ID] [Switch = none]
int_commits	The number of commit SQL statements initiated internally by the database manager during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
int_dl_rbs	The number of rollback SQL statements due to deadlocks initiated internally by the database manager during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
int_rebinds	The number of rebinds initiated internally by the database manager during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
int_rollbacks	The number of rollback SQL statements initiated internally by the database manager during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
int_rows_deleted	The number of row deletions initiated internally by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
int_rows_inserted	The number of row inserts initiated internally by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
int_rows_updated	The number of row updates initiated internally by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID] [Switch = none]
l_data_rds	The number of logical read requests of data pages by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]

l_idx_rds	The number of logical read requests of index pages by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
l_rds	The number of logical read requests of data and index pages by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
lckwtime	The amount of time in milliseconds spent by the application waiting for locks during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = lock]
loc_cursors	The number of local cursors currently open by the application. This value includes cursors that use I/O blocking, as well as those that do not. [Sequential = LST Non-Sequential = SUM] [Switch = none]
loc_cursors_blk	The number of local cursors that use I/O blocking currently open by the application [Sequential = LST Non-Sequential = SUM] [Switch = none]
lock_escals	The number of times locks by the application were escalated during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
lock_exescals	The number of times locks by the application were escalated to exclusive during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
lock_timeouts	The number of times a lock request by the application timed-out during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
lock_waits	The number of times the application had to wait for a lock during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = lock]
locks_held	The current number of locks held by the application [Sequential = LST Non-Sequential = SUM] [Switch = none]
locks_waiting	The current number of agents waiting on a lock [Sequential = LST Non-Sequential = SUM] [Switch = none]
max_assoc_agents	The maximum number of subagents associated with the application [Sequential = MAX Non-Sequential = MAX] [Switch = none]

node_num	The node identifier where the application connected to the instance [Sequential = ID Non-Sequential = ID] [Switch = none]
num_uow	The number of units of work (such as commits, rollbacks, internal commits, and internal rollbacks) generated by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
p_data_rds	The number of physical read requests of data pages by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
p_idx_rds	The number of physical read requests of index pages by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
p_rds	The number of physical read requests of data and index pages by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
p_wrts	The number of data and index pages written by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
pkgc_inserts	The number of times the application requested a section that was not available and the section had to be loaded into the package cache during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
pkgc_lookups	The number of times the application looked for a section or package in the package cache during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
pre_wait_time	The time in milliseconds spent waiting for an I/O server (prefetcher) to finish loading pages into a bufferpool in the database for the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
priority	The priority of the agents working for this application [Sequential = LST Non-Sequential = ID] [Switch = none]

priority_type	<p>The operating system priority type for the agent working on behalf of the application, represented by one of the following one-character codes:</p> <p style="margin-left: 40px;">D = Dynamic S = Static</p> <p>[Sequential = ID Non-Sequential = ID] [Switch = none]</p>
rdtime	<p>The amount of time in milliseconds spent reading pages by the application during the interval. Not stored by default.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]</p>
rej_cur_reqs	<p>The number of times that a request by the application for an I/O block was rejected during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
rem_cursors	<p>The number of remote cursors currently open by the application. This value includes cursors that use I/O blocking, as well as those that do not.</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = none]</p>
rem_cursors_blk	<p>The number of remote cursors that use I/O blocking currently open by the application</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = none]</p>
rollback_sqls	<p>The number of rollback SQL statements attempted by the application during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
rows_deleted	<p>The number of rows deleted by the application during the interval. Not stored by default.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
rows_inserted	<p>The number of rows inserted by the application during the interval. Not stored by default.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
rows_read	<p>The number of rows read from tables by the application during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
rows_selected	<p>The number of rows selected by the application during the interval. Not stored by default.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>

rows_updated	The number of rows updated by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
rows_written	The number of rows written to tables by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
Sample_End_Time	The actual timestamp of when data collection for the sample completed [Sequential = LST Non-Sequential = ID] [Switch = none]
sec_inserts	The number of inserts of SQL sections by the application from its SQL work area during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
sec_lookups	The number of lookups of SQL sections by the application from its SQL work area during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]
select_sqls	The number of SQL select statements executed by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
sort_overflows	The number of sorts by the application that ran out of sort heap and may have required disk space for temporary storage during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = sort]
sort_time	The amount of time in milliseconds spent sorting by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = sort]
sorts	The number of sorts that have been executed by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = sort]
stat_chg_time	The timestamp when the application entered its current status [Sequential = LST Non-Sequential = ID] [Switch = uow]
static_sqls	The number of static SQL statements attempted by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]

stmt_cost_est	<p>The estimated cost in timerons for the current or most recently used statement for the application. A timeron is an abstract unit of measure. It does not directly equate to any actual elapsed time, but gives an estimate of the resources (cost) required by the database manager to execute an access plan. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = statement]</p>
stmt_creator	<p>The authorization identifier of the user that pre-compiled the current or most recently used package for the application. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = statement]</p>
stmt_cursor	<p>The name of the cursor corresponding to the current or most recently used statement for the application. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = statement]</p>
stmt_etime	<p>The amount of elapsed time in seconds spent by the current or most recently used statement for the application. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = statement]</p>
stmt_operation	<p>The current or most recently used statement operation for the application. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
stmt_package	<p>The name of the current or most recently used package for the application. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = statement]</p>
stmt_row_est	<p>The estimated number of rows to be returned for the current or most recently used statement for the application. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = statement]</p>
stmt_rows_read	<p>The number of rows read from tables by the current or most recently used statement for the application. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = none]</p>
stmt_rows_written	<p>The number of rows written to tables by the current or most recently used statement for the application. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = none]</p>
stmt_sec_num	<p>The current or most recently used internal section number for the application. This value is relative to the package being used. Not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = statement]</p>

DB2 Universal Database (UDB) Server

stmt_sql_text	<p>The text of the current or most recently used statement for the application. Not stored by default. [Sequential = LST Non-Sequential = ID] [Switch = statement]</p>
stmt_sys_cpu	<p>The amount of time in seconds spent by the current or most recently used statement for the application executing system calls. Not stored by default. [Sequential = LST Non-Sequential = SUM] [Switch = statement]</p>
stmt_type	<p>The current or most recent type of statement for the application, represented by one of the following one-character codes:</p> <p style="margin-left: 40px;">D = Dynamic N = Non-SQL Statement S = Static U = Unknown</p> <p>Not stored by default. [Sequential = LST Non-Sequential = ID] [Switch = statement]</p>
stmt_usr_cpu	<p>The amount of time in seconds spent by the current or most recently used statement for the application executing database manager code. Not stored by default. [Sequential = LST Non-Sequential = SUM] [Switch = statement]</p>
System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID] [Switch = none]</p>
Time	<p>The timestamp of the data sample [Sequential = LST Non-Sequential = ID] [Switch = none]</p>
total_sqls	<p>The total number of SQL statements attempted by the application during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
tpc_accstr	<p>Data passed to the target database for logging purposes by the transaction manager. Not stored by default. [Sequential = ID Non-Sequential = ID] [Switch = none]</p>
tpc_applname	<p>The name of the application as defined by the transaction manager. Not stored by default. [Sequential = ID Non-Sequential = ID] [Switch = none]</p>
tpc_userid	<p>The client user id generated by a transaction manager. Not stored by default. [Sequential = ID Non-Sequential = ID] [Switch = none]</p>

tpc_wsname	The name of the system on which the client application is located, as defined by the transaction manager. Not stored by default. [Sequential = ID Non-Sequential = ID] [Switch = none]
uid_sqls	The number of insert, update, or delete SQL statements executed by the application during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = none]
unrd_pre_pgs	The number of prefetched pages that were unread during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
uow_comp_status	The status of the current or most recent unit of work for the application. Not stored by default. [Sequential = LST Non-Sequential = ID] [Switch = uow]
uow_etime	The amount of elapsed time in seconds for the most recently completed unit of work. Not stored by default. [Sequential = LST Non-Sequential = SUM] [Switch = uow]
uow_lckwtime	The amount of time in milliseconds spent waiting for locks by the current or most recent unit of work for the application. Not stored by default. [Sequential = LST Non-Sequential = SUM] [Switch = uow]
uow_log_space	The amount of log space in kilobytes used by the current or most recent unit of work for the application. Not stored by default. [Sequential = LST Non-Sequential = SUM] [Switch = uow]
uow_prev_stop	The timestamp of when the second most recent unit of work for the application completed. Not stored by default. [Sequential = LST Non-Sequential = ID] [Switch = uow]
uow_start_time	The timestamp of when the current or most recent unit of work for the application first requested database resources. Not stored by default. [Sequential = LST Non-Sequential = ID] [Switch = uow]
uow_stop_time	The timestamp of when the most recent unit of work for the application completed. Not stored by default. [Sequential = LST Non-Sequential = ID] [Switch = uow]
wks_p_inserts	The number of times an insert of SQL sections by the application occurred in private workspaces during the interval. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>. [Sequential = SUM Non-Sequential = SUM] [Switch = none]

wks_p_lookups	<p>The number of times a lookup of SQL sections by the application occurred in private workspaces during the interval. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
wks_p_max_mem	<p>The maximum amount of memory in kilobytes used in private workspaces by this application. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = MAX Non-Sequential = SUM] [Switch = none]</p>
wks_p_overflows	<p>The number of times private workspaces overflowed the bounds of their allocated memory due to the application during the interval. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
wks_s_inserts	<p>The number of times an insert of SQL sections by the application occurred in shared workspaces during the interval. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
wks_s_lookups	<p>The number of times a lookup of SQL sections by the application occurred in shared workspaces during the interval. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
wks_s_max_mem	<p>The maximum amount of memory in kilobytes used in shared workspaces by this application. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = MAX Non-Sequential = SUM] [Switch = none]</p>
wks_s_overflows	<p>The number of times shared workspaces overflowed the bounds of their allocated memory due to the application during the interval. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
wrttime	<p>The amount of time in milliseconds spent writing pages by the application during the interval. Not stored by default.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]</p>

3.2. Bufferpool Statistics

Detailed information about the bufferpools within each DB2 database is stored within the DB2.Bufferpool Detail table in the TeamQuest performance database. To obtain this information, the bufferpool DB2 monitor switch must be enabled within the DB2 UDB Agent configuration entry for the instance.

Table Field Hierarchy

Class:	DB2
Subclass:	Bufferpool Detail
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Bufferpool Detail
Open Table Name:	DB2BUFFPOOLDETAIL
Collection interval:	1 minute (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
a_data_rd_reqs	The number of asynchronous data read requests from the bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_data_rds	The number of data pages read asynchronously into the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_data_wrts	The number of data pages written to disk asynchronously from the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_idx_rd_reqs	The number of asynchronous index read requests from the bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_idx_rds	The number of index pages read asynchronously into the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_idx_wrts	The number of index pages written to disk asynchronously from the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]

a_rds	The number of data and index pages read asynchronously into the bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_rdtype	The amount of time in milliseconds spent reading pages asynchronously into the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_wrts	The number of data and index pages written to disk asynchronously from the bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_wrttime	The amount of time in milliseconds spent writing pages to disk asynchronously from the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID] [Switch = none]
avg_a_data_rds	The average number of data pages read per asynchronous data read request during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_a_idx_rds	The average number of index pages read per asynchronous index read request during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_a_rdtype	The average time in milliseconds spent reading a page asynchronously into the bufferpool during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_a_wrttime	The average time in milliseconds spent writing a page to disk asynchronously from the bufferpool during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_rds	The average number of sectors that are read per direct read request during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_rdtype	The average time in milliseconds spent direct reading a sector during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]

avg_d_wrts	The average number of sectors that are written per direct write request during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_wrttime	The average time in milliseconds spent direct writing a sector during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_rdttime	The average time in milliseconds spent reading a page into the bufferpool during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_wrttime	The average time in milliseconds spent writing a page to disk from the bufferpool during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
block_pgs	The number of pages read by block I/O during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
block_reqs	The number of block I/O requests made during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
bufferpool_name	The name of the bufferpool [Sequential = ID Non-Sequential = ID] [Switch = none]
curr_size	The current size in pages of the bufferpool [Sequential = LST Non-Sequential = SUM] [Switch = none]
d_rd_reqs	The number of requests to perform direct read operations during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_rds	The number of direct read operations during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_rdttime	The amount of time in milliseconds spent direct reading a sector during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_wrt_reqs	The number of requests to perform direct write operations during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_wrts	The number of direct write operations during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]

d_wrttime	The amount of time in milliseconds spent direct writing a sector during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
data_wrts	The number of data pages written to disk from the bufferpool during the interval. This includes writes done asynchronously. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
database	The name of the database [Sequential = ID Non-Sequential = ID] [Switch = none]
files_closed	The number of times a file had to be closed to remain under the maxfilop configuration parameter when opening a new file for use with the bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
hit_ratio	The percentage of logical reads that were satisfied from the bufferpool without having to invoke a read from disk during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
idx_wrts	The number of index pages written to disk from the bufferpool during the interval. This includes writes done asynchronously. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
Instance	The name of the instance from which data is obtained. Up to 20 characters are displayed. [Sequential = ID Non-Sequential = ID] [Switch = none]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID] [Switch = none]
l_data_rds	The number of logical read requests of data pages from the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
l_idx_rds	The number of logical read requests of index pages from the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
l_rds	The number of logical read requests of data and index pages from the bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]

new_size	The size in pages the bufferpool will be changed to once the database is restarted [Sequential = LST Non-Sequential = SUM] [Switch = none]
node_num	The node number where this bufferpool resides [Sequential = LST Non-Sequential = ID] [Switch = none]
num_ts	The number of tablespaces using the bufferpool [Sequential = LST Non-Sequential = SUM] [Switch = none]
p_data_rds	The number of physical read requests of data pages from the bufferpool during the interval. This includes reads done asynchronously. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
p_idx_rds	The number of physical read requests of index pages from the bufferpool during the interval. This includes reads done asynchronously. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
p_rds	The number of physical read requests of data and index pages from the bufferpool during the interval. This includes reads done asynchronously. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
p_wrts	The number of data and index pages written to disk from the bufferpool during the interval. This includes writes done asynchronously. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
phy_pg_maps	The number of physical page maps made during the interval. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
rdtime	The amount of time in milliseconds spent reading pages into the bufferpool during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
remove_pgs_left	The number of pages left to remove from the bufferpool before the bufferpool resize is completed [Sequential = LST Non-Sequential = SUM] [Switch = none]
Sample_End_Time	The actual timestamp of when data collection for the sample completed [Sequential = LST Non-Sequential = ID] [Switch = none]

DB2 Universal Database (UDB) Server

System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p> <p>[Switch = none]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p> <p>[Switch = none]</p>
unrd_pre_pgs	<p>The number of prefetched pages that were unread during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
vect_pgs	<p>The number of pages read by vectored I/O during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
vect_reqs	<p>The number of vectored I/O requests made during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
wrttime	<p>The amount of time in milliseconds spent writing pages to disk from the bufferpool during the interval. Not stored by default.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>

3.3. Database Configuration Statistics

Detailed information about the configuration of DB2 databases is stored within the DB2.Database Config table in the TeamQuest performance database. Information is only collected for databases listed in the Database List setting of the DB2 UDB Agent configuration entry for the instance.

Table Field Hierarchy

Class:	DB2
Subclass:	Database Config
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Database Config
Open Table Name:	DB2DBCONFIG
Collection interval:	N/A
Default retention:	1 year
Table type:	Event

Statistic Name	Description
database	The name of the database [Non-Sequential = ID] [Switch = none]
db_alias	The alias used to refer to the database [Non-Sequential = ID] [Switch = none]
Instance	The name of the instance from which data is obtained. Up to 24 characters are displayed. [Non-Sequential = ID] [Switch = none]
parameter	The name of the database manager configuration parameter [Non-Sequential = ID] [Switch = none]
Sample_End_Time	The actual timestamp of when data collection for the sample completed [Non-Sequential = ID] [Switch = none]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID] [Switch = none]

Time	The timestamp of the data sample [Non-Sequential = ID] [Switch = none]
value	The value of the database manager configuration parameter [Non-Sequential = ID] [Switch = none]

3.4. Database Statistics

Performance data relating to each specific DB2 database is stored as parameters in the TeamQuest performance database. The statistics are classified by the hierarchy of key names. A statistic marked with an asterisk (*) is a derived statistic.

Note: *Resource names for database statistics are a concatenation of the instance name and the database name, separated by a period (for example, db2stat1.prod).*

Parameter Hierarchy

Class:	DB2
Subclass:	Database.Agent
IT Resource Name:	/TeamQuest/System/systemname/DB2
TeamQuest Table Name:	DB2.Database.Agent
Open Table Name:	DB2DBAGENT
Resource:	database1, database2, ...
Statistic Name:	
CurrentAgents	The current number of subagents for all applications within the database [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/AgentCounts.rpt
MaxAgents	The maximum number of agents for all applications since the database was activated [Sequential = MAX Non-Sequential = SUM] [Switch = statement] View Report: /report/db2/database/AgentCounts.rpt
MaxCoordAgents	The maximum number of coordinating agents working at one time since the database was activated [Sequential = MAX Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/AgentCounts.rpt

Class:	DB2
Subclass:	Database.Bufferpool Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Bufferpool Summary
Open Table Name:	DB2DBBUFFPOOLSUM
Resource:	database1, database2, ...
Statistic Name:	
AsyncDataReads/s	The number of data pages per second read asynchronously into all bufferpools within the database [Sequential = AVG Non-Sequential = SUM] [Switch = bufferpool]
AsyncDataWrites/s	The number of data pages per second written to disk asynchronously from all bufferpools within the database [Sequential = AVG Non-Sequential = SUM] [Switch = bufferpool]
AsyncIndexReads/s	The number of index pages per second read asynchronously into all bufferpools within the database [Sequential = AVG Non-Sequential = SUM] [Switch = bufferpool]
AsyncIndexWrites/s	The number of index pages per second written to disk asynchronously from all bufferpools within the database [Sequential = AVG Non-Sequential = SUM] [Switch = bufferpool]
AsyncReads/s*	The number of data and index pages per second read asynchronously into all bufferpools within the database [Switch = bufferpool] View Report: /report/db2/database/BufferpoolIORates.rpt
AsyncWrites/s*	The number of data and index pages per second written to disk asynchronously from all bufferpools within the database [Switch = bufferpool] View Report: /report/db2/database/BufferpoolIORates.rpt
AvgAsyncDataReads	The average number of data pages read per asynchronous data read request [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
AvgAsyncIndexReads	The average number of index pages read per asynchronous index read request [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]

AvgAsyncReadTime	The average time in milliseconds spent reading a page asynchronously into a bufferpool in the database [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool] View Report: /report/db2/database/BufferpoolTime.rpt
AvgAsyncWriteTime	The average time in milliseconds spent writing a page asynchronously from a bufferpool in the database to disk [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool] View Report: /report/db2/database/BufferpoolTime.rpt
AvgDirectReads	The average number of sectors that are read per direct read request [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
AvgDirectReadTime	The average time in milliseconds spent direct reading a sector [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool] View Report: /report/db2/database/BufferpoolTime.rpt
AvgDirectWrites	The average number of sectors that are written per direct write request [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
AvgDirectWriteTime	The average time in milliseconds spend direct writing a sector [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool] View Report: /report/db2/database/BufferpoolTime.rpt
AvgReadTime	The average time in milliseconds spent reading a page into a bufferpool in the database [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool] View Report: /report/db2/database/BufferpoolTime.rpt
AvgWriteTime	The average time in milliseconds spent writing a page from a bufferpool in the database to disk [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool] View Report: /report/db2/database/BufferpoolTime.rpt
DataWrites/s	The number of data pages per second written to disk from all bufferpools within the database. This includes writes done asynchronously. [Sequential = AVG Non-Sequential = SUM] [Switch = bufferpool]

DBFilesClosed/s	<p>The number of times per second a file had to be closed to remain under the maxfilop configuration parameter when opening a new file for use with a bufferpool in the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p> <p>View Report: /report/db2/database/BufferpoolPagingRates.rpt</p>
DirectReads/s	<p>The number of read operations per second that do not use a bufferpool in the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p> <p>View Report: /report/db2/database/BufferpoolIORates.rpt</p>
DirectWrites/s	<p>The number of write operations per second that do not use a bufferpool in the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p> <p>View Report: /report/db2/database/BufferpoolIORates.rpt</p>
HitRatioPct*	<p>The percentage of logical reads that were satisfied from the bufferpools within the database without having to invoke a read from disk</p> <p>[Switch = bufferpool]</p> <p>View Report: /report/db2/database/MemoryHitRatios.rpt</p>
IndexWrites/s	<p>The number of index pages per second written to disk from all bufferpools within the database. This includes writes done asynchronously.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
LogicalDataReads/s	<p>The number of logical read requests per second of data pages from all the bufferpools within the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
LogicalIndexReads/s	<p>The number of logical read requests per second of index pages from all bufferpools within the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
LogicalReads/s*	<p>The number of logical read requests per second of data and index pages from all bufferpools within the database</p> <p>[Switch = bufferpool]</p>
LSNCleans/s	<p>The number of times per second a page cleaner was invoked because the logging space used had reached a predefined criterion for the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p> <p>View Report: /report/db2/database/BufferpoolPagingRates.rpt</p>

PhysicalDataReads/s	<p>The number of physical read requests per second of data pages from all the bufferpools within the database. This includes reads done asynchronously.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
PhysicalIndexReads/s	<p>The number of physical read requests per second of index pages from all bufferpools within the database. This includes reads done asynchronously.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
PhysicalReads/s*	<p>The number of physical read requests per second of data and index pages from all bufferpools within the database</p> <p>[Switch = bufferpool]</p> <p>View Report:</p> <p>/report/db2/database/BufferpoolIORates.rpt</p>
PhysicalWrites/s*	<p>The number of data and index pages per second written to disk from all bufferpools within the database</p> <p>[Switch = bufferpool]</p> <p>View Report:</p> <p>/report/db2/database/BufferpoolIORates.rpt</p>
PrefetchWaitTime	<p>The time in milliseconds spent by applications waiting for an I/O server (prefetcher) to finish loading pages into a bufferpool in the database</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
ThresholdCleans/s	<p>The number of times per second a page cleaner was invoked because a bufferpool had reached the dirty page threshold criteria for the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p> <p>View Report:</p> <p>/report/db2/database/BufferpoolPagingRates.rpt</p>
UnreadPrefetch Pages/s	<p>The number of prefetch pages per second that were unread</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p> <p>View Report:</p> <p>/report/db2/database/BufferpoolPagingRates.rpt</p>
VictimPageCleans/s	<p>The number of times per second a page cleaner was invoked because a synchronous write was needed during the victim replacement for the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p> <p>View Report:</p> <p>/report/db2/database/BufferpoolPagingRates.rpt</p>

Class:	DB2
Subclass:	Database.Catalog Cache
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Catalog Cache
Open Table Name:	DB2DBCATALOGCACHE
Resource:	database1, database2, ...
Statistic Name:	
HeapFullOverflows/s	<p>The number of times per second that an insert into the catalog cache failed due to a heap-full condition in the database heap [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/CacheOverflows.rpt</p>
HitRatioPct*	<p>The percentage of lookups satisfied from the catalog cache without incurring an insert [Switch = none] View Report: /report/db2/database/MemoryHitRatios.rpt</p>
Inserts/s	<p>The number of times per second the system tried to insert table descriptor information into the catalog cache within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]</p>
Lookups/s	<p>The number of times per second the catalog cache was referenced to obtain table descriptor information within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]</p>
MaxMemoryUsed	<p>The maximum amount of memory in megabytes used by the catalog cache since the database was activated [Sequential = MAX Non-Sequential = SUM] [Switch = none]</p>
Overflows/s	<p>The number of times per second that an insert into the catalog cache failed due to the catalog cache being full [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/CacheOverflows.rpt</p>

DB2 Universal Database (UDB) Server

Class:	DB2
Subclass:	Database.Connection
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Connection
Open Table Name:	DB2DBCONN
Resource:	database1, database2, ...
Statistic Name:	
AvgConnections	The average number of applications connected to the database over the interval [Sequential = AVG Non-Sequential = SUM] [Switch = none]
Connections/s	The number of connections per second made to the database [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/ConnectionRates.rpt
CurrentConnections	The current number of applications connected to the database [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/ConnectionCounts.rpt
Executing	The current number of applications connected to the database that are processing a unit of work [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/ConnectionCounts.rpt
MaxConnections	The maximum number of simultaneous connections to the database since it was activated [Sequential = MAX Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/ConnectionCounts.rpt
SecConnections/s	The number of connections per second made by a subagent to the database [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/ConnectionRates.rpt

Class:	DB2
Subclass:	Database.Hash Join
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Hash Join
Open Table Name:	DB2DBHASHJOIN
Resource:	database1, database2, ...
Statistic Name:	
HashJoins/s	The number of hash joins executed per second within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
Loops/s	The number of times per second a single partition hash join was larger than the available sort heap space in the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
Overflows/s	The number of times per second that hash join data exceeded the available sort heap space [Sequential = AVG Non-Sequential = SUM] [Switch = none]
SmallOverflows/s	The number of times per second that hash join data exceeded the available sort heap space by less than 10% [Sequential = AVG Non-Sequential = SUM] [Switch = none]
Class:	DB2
Subclass:	Database.Lock
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Lock
Open Table Name:	DB2DBLK
Resource:	database1, database2, ...
Statistic Name:	
AvgWaitTime	The average amount of time in milliseconds spent by applications within the database waiting for a lock [Sequential = AVG Non-Sequential = AVG] [Switch = lock] View Report: /report/db2/database/LockTime.rpt
Deadlocks/s	The number of deadlocks per second within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/LockRates.rpt

Escalations/s	<p>The number of times per second locks were escalated within the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LockRates.rpt</p>
ExEscalations/s	<p>The number of times per second locks were escalated to exclusive within the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LockRates.rpt</p>
Held	<p>The current number of locks held by all applications in the database</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LockCounts.rpt</p>
MemoryUsed	<p>The current amount of lock list memory in megabytes in use within the database</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/MemoryUsage.rpt</p>
Timeouts/s	<p>The number of times per second a lock request timed out within the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LockRates.rpt</p>
Waiting	<p>The current number of agents waiting on a lock within the database</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LockCounts.rpt</p>
Waits/s	<p>The number of times per second applications within the database had to wait for a lock</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = lock]</p> <p>View Report:</p> <p>/report/db2/database/LockRates.rpt</p>
WaitTime	<p>The amount of time in milliseconds spent by applications within the database waiting for locks</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>[Switch = lock]</p>

Class:	DB2
Subclass:	Database.Log Space
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Log Space
Open Table Name:	DB2DBLOGSPACE
Resource:	database1, database2, ...
Statistic Name:	
Available	<p>The current amount of log space in megabytes available within the database</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LogSpaceUsage.rpt</p>
IndoubtTransactions	<p>The number of outstanding indoubt transactions in the database</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p>
MaxSecUsed	<p>The maximum amount of secondary log space in megabytes used since the database was activated</p> <p>[Sequential = MAX Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LogSpaceUsage.rpt</p>
MaxUsed	<p>The maximum amount of log space in megabytes used since the database was activated</p> <p>[Sequential = MAX Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LogSpaceUsage.rpt</p>
Reads/s	<p>The number of log page reads per second within the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LogSpaceRates.rpt</p>
SecLogsAllocated	<p>The current number of secondary log files that are being used for the database</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p>
Used	<p>The current amount of log space in megabytes used within the database</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LogSpaceUsage.rpt</p>
Writes/s	<p>The number of log page writes per second within the database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/database/LogSpaceRates.rpt</p>

DB2 Universal Database (UDB) Server

Class:	DB2
Subclass:	Database.Package Cache
IT Resource Name:	/TeamQuest/System/systemname/DB2
TeamQuest Table Name:	DB2.Database.Package Cache
Open Table Name:	DB2DBPACKAGECACHE
Resource:	database1, database2, ...
Statistic Name:	
HitRatioPct*	The percentage of lookups satisfied from the package cache without incurring an insert [Switch = none] View Report: /report/db2/database/MemoryHitRatios.rpt
Inserts/s	The number of times per second a requested section was not available and had to be loaded into the package cache within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
Lookups/s	The number of times per second an application looked for a section or package in the package cache within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
MaxMemoryUsed	The maximum amount of memory in megabytes used by the package cache since the database was activated [Sequential = MAX Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/PackageCacheMaxMemory.rpt
Overflows/s	The number of times per second that the package cache overflowed the bounds of its allocated memory within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/CacheOverflows.rpt
SectionInserts/s	The number of inserts of SQL sections per second by an application from its SQL work area within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
SectionLookups/s	The number of lookups of SQL sections per second by an application from its SQL work area within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]

Class:	DB2
Subclass:	Database.Sample
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Sample
Open Table Name:	DB2DBSAMPLE
Resource:	database1, database2, ...
Statistic Name:	
Etime	<p>The amount of time in seconds, elapsed between two samples for this DB2 database. This value may differ from the DB2.Sample..tqdb2p_interval statistic when a database is activated after data collection has begun for this instance of the DB2 UDB Agent. In this case, the value is the difference between the database activation timestamp and the current snapshot timestamp.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = none]</p>
Class:	DB2
Subclass:	Database.Sort
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Sort
Open Table Name:	DB2DBSORT
Resource:	database1, database2, ...
Statistic Name:	
AvgSortTime	<p>The average amount of time in milliseconds per sort within the database</p> <p>[Sequential = AVG Non-Sequential = AVG] [Switch = sort] View Report: /report/db2/database/SortTime.rpt</p>
CurrentSorts	<p>The current number of sorts in the database that have sort heap allocated</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/SortCounts.rpt</p>
MaxSharedMemoryUsed	<p>The maximum amount of sort heap memory allocated in megabytes for sorts at one time since the database was activated</p> <p>[Sequential = MAX Non-Sequential = SUM] [Switch = none]</p>

DB2 Universal Database (UDB) Server

MemoryUsed	The current amount of sort heap memory allocated in megabytes for all sorts in the database [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/MemoryUsage.rpt
Overflows/s	The number of sorts per second that ran out of sort heap and may have required disk space for temporary storage within the database [Sequential = AVG Non-Sequential = SUM] [Switch = sort] View Report: /report/db2/database/SortRates.rpt
SharedMemoryUsed	The current amount of shared sort heap memory allocated in megabytes for all sorts in the database [Sequential = LST Non-Sequential = SUM] [Switch = none]
Sorts/s	The number of sorts per second that have been executed within the database [Sequential = AVG Non-Sequential = SUM] [Switch = sort] View Report: /report/db2/database/SortRates.rpt
SortTime	The amount of time in milliseconds spent sorting within the database [Sequential = SUM Non-Sequential = SUM] [Switch = sort]
Class:	DB2
Subclass:	Database.Statement Summary
IT Resource Name:	/TeamQuest/System/systemname/DB2
TeamQuest Table Name:	DB2.Database.Statement Summary
Open Table Name:	DB2DBSTMTSUM
Resource:	database1, database2, ...
Statistic Name:	
Binds/s	The number of binds and pre-compiles per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
CommitSQLs/s	The number of commit SQL statements per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]

CompletedSQLs/s*	The number of dynamic and static SQL statements, minus the number of failed SQL statements, that occurred per second within the database [Switch = none] View Report: /report/db2/database/SQLRates.rpt
DDLSQLs/s	The number of data definition language (DDL) SQL statements per second executed within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
DynamicSQLs/s	The number of dynamic SQL statements per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
ExecutionTime	The amount of time in milliseconds spent executing statements within the database [Sequential = LST Non-Sequential = SUM] [Switch = statement]
FailedSQLs/s	The number of failed SQL statements per second within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
IntCommitSQLs/s	The number of commit SQL statements per second initiated internally by the database manager within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
IntDLRollbackSQLs/s	The number of rollback SQL statements per second due to deadlocks initiated internally by the database manager within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
IntRebinds/s	The number of rebinds per second initiated internally by the database manager within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
IntRollbackSQLs/s	The number of rollback SQL statements per second initiated internally by the database manager within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
RollbackSQLs/s	The number of rollback SQL statements per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
SelectSQLs/s	The number of select SQL statements per second executed within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]

DB2 Universal Database (UDB) Server

StaticSQLs/s	The number of static SQL statements per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
TotalSQLs/s*	The number of dynamic and static SQL statements per second attempted within the database [Switch = none] View Report: /report/db2/database/SQLRates.rpt
UIDSQLs/s	The number of update, insert, or delete SQL statements per second executed within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
UnitsOfWork/s*	The number of units of work per second (such as commits, rollbacks, internal commits, and internal rollbacks) generated within the database [Switch = none] View Report: /report/db2/database/SQLRates.rpt
Class:	DB2
Subclass:	Database.Table Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2
TeamQuest Table Name:	DB2.Database.Table Summary
Open Table Name:	DB2DBTABSUM
Resource:	database1, database2, ...
Statistic Name:	
IntRowsDeleted/s	The number of row deletions per second initiated internally within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
IntRowsInserted/s	The number of row inserts per second initiated internally within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
IntRowsUID/s*	The number of rows updated, inserted, and deleted per second that were initiated internally within the database [Switch = none]
IntRowsUpdated/s	The number of row updates per second initiated internally within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
RowsDeleted/s	The number of row deletions per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]

RowsInserted/s	The number of row insertions per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
RowsRead/s	The number of row reads per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/database/TableRates.rpt
RowsSelected/s	The number of row selects per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
RowsUID/s*	The number of rows updated, inserted, and deleted per second within the database [Switch = none] View Report: /report/db2/database/TableRates.rpt
RowsUpdated/s	The number of row updates per second attempted within the database [Sequential = AVG Non-Sequential = SUM] [Switch = none]
Class:	DB2
Subclass:	Database.Workspace
IT Resource Name:	/TeamQuest/System/systemname/DB2
TeamQuest Table Name:	DB2.Database.Workspace
Open Table Name:	DB2DBWORKSPACE
Resource:	database1, database2, ...
Statistic Name:	
MaxPrivateMemoryUsed	The maximum amount of memory in megabytes used at one time by private workspaces in the database. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>. [Sequential = MAX Non-Sequential = SUM] [Switch = none]
MaxSharedMemoryUsed	The maximum amount of memory in megabytes used at one time by shared workspaces in the database. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>. [Sequential = MAX Non-Sequential = SUM] [Switch = none]

PrivateInserts/s	<p>The number of times per second an insert of SQL sections by applications in the database occurred in private workspaces. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p>
PrivateLookups/s	<p>The number of times per second a lookup of SQL sections by applications in the database occurred in private workspaces. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p>
PrivateOverflows/s	<p>The number of times per second private workspaces in the database overflowed the bounds of their allocated memory. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p>
SharedInserts/s	<p>The number of times per second an insert of SQL sections by applications in the database occurred in shared workspaces. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p>
SharedLookups/s	<p>The number of times per second a lookup of SQL sections by applications in the database occurred in shared workspaces. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p>
SharedOverflows/s	<p>The number of times per second shared workspaces in the database overflowed the bounds of their allocated memory. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p>

3.5. Database Status Statistics

Detailed information about the status of DB2 databases is stored within the DB2.Database Status table in the TeamQuest performance database.

Table Field Hierarchy

Class:	DB2
Subclass:	Database Status
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Database Status
Open Table Name:	DB2DBSTATUS
Collection interval:	1 minute (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = ID Non-Sequential = ID] [Switch = none]
backup_pend	Indicates whether a full backup must be performed on the database before accessing it. This field will only have a value if the appropriate database alias has been specified in the Database List setting of the DB2 UDB Agent configuration entry for the instance. [Sequential = LST Non-Sequential = ID] [Switch = none]
cat_node	The node identifier where the database catalog tables are stored [Sequential = LST Non-Sequential = ID] [Switch = none]
cat_node_name	The network name of the catalog node [Sequential = LST Non-Sequential = ID] [Switch = none]
database	The name of the database [Sequential = ID Non-Sequential = ID] [Switch = none]
database_path	The full path of the location where the database is stored [Sequential = LST Non-Sequential = ID] [Switch = none]
db_alias	The alias used to refer to the database [Sequential = LST Non-Sequential = ID] [Switch = none]

DB2 Universal Database (UDB) Server

db_consist	<p>Indicates whether the database is in a consistent state. This field will only have a value if the appropriate database alias has been specified in the Database List setting of the DB2 UDB Agent configuration entry for the instance.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
db_loc	<p>The location of the database in relation to the DB2 UDB Agent, represented by one of the following one-character codes:</p> <p style="padding-left: 40px;">L = Local R = Remote</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
Instance	<p>The name of the instance from which data is obtained. Up to 24 characters are displayed.</p> <p>[Sequential = ID Non-Sequential = ID] [Switch = none]</p>
Interval	<p>The expected sampling interval in seconds</p> <p>[Sequential = SUM Non-Sequential = ID] [Switch = none]</p>
last_backup	<p>The timestamp of when the last backup occurred for the database. A value of <N/A> will be stored if no backups have occurred.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
log_head	<p>The name of the log file that is currently active. This field will only have a value if the appropriate database alias has been specified in the Database List setting of the DB2 UDB Agent configuration entry for the instance.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
log_path	<p>The current path being used for logging purposes. This field will only have a value if the appropriate database alias has been specified in the Database List setting of the DB2 UDB Agent configuration entry for the instance.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
restore_pend	<p>Indicates whether a restore pending status exists in the database. This field will only have a value if the appropriate database alias has been specified in the Database List setting of the DB2 UDB Agent configuration entry for the instance.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>

roll_pend	Indicates whether a database or tablespace rollforward is needed. Otherwise, it will indicate that no rollforward recovery is required. This field will only have a value if the appropriate database alias has been specified in the Database List setting of the DB2 UDB Agent configuration entry for the instance. [Sequential = LST Non-Sequential = ID] [Switch = none]
Sample_End_Time	The actual timestamp of when data collection for the sample completed [Sequential = LST Non-Sequential = ID] [Switch = none]
svr_platform	The operating system on which the database server is running [Sequential = LST Non-Sequential = ID] [Switch = none]
start_time	The timestamp of when the database was activated [Sequential = LST Non-Sequential = ID] [Switch = none]
status	The current status of the database. A value is always present while the database is active. If the database is inactive and a valid database alias has been specified in the Database List setting of the DB2 UDB Agent configuration entry for the instance, a value of Inactive is stored. [Sequential = LST Non-Sequential = ID] [Switch = none]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID] [Switch = none]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID] [Switch = none]

3.6. Instance Configuration Statistics

Detailed information about the configuration of the DB2 instance is stored within the DB2.Instance Config table in the TeamQuest performance database.

Table Field Hierarchy

Class:	DB2
Subclass:	Instance Config
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Instance Config
Open Table Name:	DB2INSTANCECONFIG
Collection interval:	N/A
Default retention:	1 year
Table type:	Event

Statistic Name	Description
Instance	The name of the instance from which data is obtained. Up to 24 characters are displayed. [Non-Sequential = ID] [Switch = none]
parameter	The name of the database manager configuration parameter [Non-Sequential = ID] [Switch = none]
Sample_End_Time	The actual timestamp of when data collection for the sample completed [Non-Sequential = ID] [Switch = none]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID] [Switch = none]
Time	The timestamp of the data sample [Non-Sequential = ID] [Switch = none]
value	The value of the database manager configuration parameter [Non-Sequential = ID] [Switch = none]

3.7. Instance Statistics

Data pertaining to the overall DB2 UDB instance is stored as parameters in the TeamQuest performance database. The statistics are classified by the hierarchy of key names. A statistic marked with an asterisk (*) is a derived statistic.

Parameter Hierarchy

Class:	DB2
Subclass:	Instance.Agent
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Instance.Agent
Open Table Name:	DB2INSTANCEAGENT
Resource:	instance1, instance2, ...
Statistic Name:	
Assigns/s	<p>The number of agents assigned per second from the agent pool [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/instance/AgentRates.rpt</p>
Creates/s	<p>The number of agents created per second because the agent pool was empty. This value includes the number of agents started when the instance began. [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/instance/AgentRates.rpt</p>
Idle	<p>The current number of agents (coordinator agents and subagents) that are unassigned to an application [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: /report/db2/instance/AgentCounts.rpt</p>
MaxCoordAgents	<p>The maximum number of coordinating agents working at one time since the instance started [Sequential = MAX Non-Sequential = SUM] [Switch = none]</p>
MaxRegistered	<p>The maximum number of agents registered at one time since the instance started [Sequential = MAX Non-Sequential = SUM] [Switch = none]</p>
MaxWaiting	<p>The maximum number of agents waiting for a token at one time since the instance started. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>. [Sequential = MAX Non-Sequential = SUM] [Switch = none]</p>

DB2 Universal Database (UDB) Server

Overflows/s	<p>The number of times per second a request to create a new agent was received after the maxagents configuration parameter had already been reached. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/AgentRates.rpt</p>
Registered	<p>The current number of agents (coordinator agents and subagents) registered for the instance</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/AgentCounts.rpt</p>
Steals/s	<p>The number of times per second an agent was stolen from an application</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/AgentRates.rpt</p>
Waiting	<p>The current number of agents (coordinator agents and subagents) waiting for a token. This statistic is not available for DB2 UDB 9.5 and later. The value is displayed as <N/A>.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/AgentCounts.rpt</p>
Class:	DB2
Subclass:	Instance.Connection
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Instance.Connection
Open Table Name:	DB2INSTANCECONN
Resource:	instance1, instance2, ...
Statistic Name:	
AvgConnections	<p>The average number of connections to the instance from local and remote clients during the interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p>
Local	<p>The current number of connections to the instance from local clients</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/ConnectionCounts.rpt</p>

LocalDBsWithCons	<p>The current number of local databases with applications connected</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p>
LocalExec	<p>The current number of connections to the instance from local clients that are processing a unit of work</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/ConnectionCounts.rpt</p>
Remote	<p>The current number of connections to the instance from remote clients</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/ConnectionCounts.rpt</p>
RemoteExec	<p>The current number of connections to the instance from remote clients that are processing a unit of work</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/ConnectionCounts.rpt</p>
Class:	DB2
Subclass:	Instance.Gateway
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2/ <i>instancename</i>
TeamQuest Table Name:	DB2.Instance.Gateway
Open Table Name:	DB2INSTANCEGATEWAY
Resource:	instance1, instance2, ...
Statistic Name:	
AgentSteals/s	<p>The number of times per second that an agent from the agents pool was primed with a connection and was stolen for use with a different Distributed Relational Database Architecture (DRDA) database</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/GatewayRates.rpt</p>
Connections/s	<p>The number of connections per second attempted from the DB2 Connect gateway</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>[Switch = none]</p> <p>View Report:</p> <p>/report/db2/instance/GatewayRates.rpt</p>

DB2 Universal Database (UDB) Server

ConsWaitingOnClient	The current number of connections to host databases being handled by the DB2 Connect gateway that are waiting for the client to send a request [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: <i>/report/db2/instance/GatewayCounts.rpt</i>
ConsWaitingOnHost	The current number of connections to host databases being handled by the DB2 Connect gateway that are waiting for a reply from the host [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: <i>/report/db2/instance/GatewayCounts.rpt</i>
CurrentConnections	The current number of connections to host databases being handled by the DB2 Connect gateway [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: <i>/report/db2/instance/GatewayCounts.rpt</i>
Class:	DB2
Subclass:	Instance.Hash Join
IT Resource Name:	<i>/TeamQuest/System/systemname/DB2/instancename</i>
TeamQuest Table Name:	DB2.Instance.Hash Join
Open Table Name:	DB2INSTANCEHASHJOIN
Resource:	instance1, instance2, ...
Statistic Name:	
Overflows/s	The number of times per second a hash join heap request was limited due to concurrent use of the shared or private sort heap space [Sequential = AVG Non-Sequential = SUM] [Switch = none]
Class:	DB2
Subclass:	Instance.Memory
IT Resource Name:	<i>/TeamQuest/System/systemname/DB2/instancename</i>
TeamQuest Table Name:	DB2.Instance.Memory
Open Table Name:	DB2INSTANCEMEM
Resource:	instance1, instance2, ...
Statistic Name:	
PrivateMemCommitted	The current amount of private memory committed in megabytes [Sequential = LST Non-Sequential = SUM] [Switch = none]

Class:	DB2
Subclass:	Instance.Sort
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2/ <i>instancename</i>
TeamQuest Table Name:	DB2.Instance.Sort
Open Table Name:	DB2INSTANCESORT
Resource:	instance1, instance2, ...
Statistic Name:	
MaxMemoryUsed	The maximum amount of sort heap memory allocated in megabytes for sorts at one time since the instance started [Sequential = MAX Non-Sequential = SUM] [Switch = none]
MemoryUsed	The current amount of sort heap memory allocated in megabytes for all sorts in the instance [Sequential = LST Non-Sequential = SUM] [Switch = none] View Report: /report/db2/instance/SortMemory.rpt
Overflows/s	The number of sorts per second that have requested heaps after the sort heap threshold has been reached [Sequential = AVG Non-Sequential = SUM] [Switch = sort] View Report: /report/db2/instance/SortRates.rpt
PipedAccepts/s	The number of piped sorts per second that have been accepted [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/instance/SortRates.rpt
PipedRequests/s	The number of piped sorts per second that have been requested [Sequential = AVG Non-Sequential = SUM] [Switch = none] View Report: /report/db2/instance/SortRates.rpt

Class:	DB2
Subclass:	Sample
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Sample
Open Table Name:	DB2SAMPLE
Resource:	instance1, instance2, ...
Statistic Name:	
tqdb2p_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = SUM] [Switch = none]
tqdb2p_interval	The amount of time in seconds elapsed between two samples of this instance of the DB2 UDB Agent [Sequential = SUM Non-Sequential = SUM] [Switch = none]

3.8. Instance Status Statistics

Detailed information about the status of the DB2 instance is stored within the DB2.Instance Status table in the TeamQuest performance database.

Table Field Hierarchy

Class:	DB2
Subclass:	Instance Status
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Instance Status
Open Table Name:	DB2INSTANCESTATUS
Collection interval:	1 minute (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID] [Switch = none]
Instance	The name of the instance from which data is obtained. Up to 24 characters are displayed. [Sequential = ID Non-Sequential = ID] [Switch = none]

Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID] [Switch = none]
node_num	The node number to which the DB2 UDB Agent is connected [Sequential = LST Non-Sequential = ID] [Switch = none]
product_name	The detailed version of the DB2 instance [Sequential = LST Non-Sequential = ID] [Switch = none]
Sample_End_Time	The actual timestamp of when data collection for the sample completed [Sequential = LST Non-Sequential = ID] [Switch = none]
service_level	The current corrective service level of the DB2 instance [Sequential = LST Non-Sequential = ID] [Switch = none]
svr_instance	The name of the instance as reported by the database manager being monitored [Sequential = LST Non-Sequential = ID] [Switch = none]
svr_node	The name of the system on which the database manager being monitored resides [Sequential = LST Non-Sequential = ID] [Switch = none]
svr_prdid	The product and version of the database manager. The string is in the form <i>SQLVRRM</i> , where <i>VV</i> is a 2-digit version number, <i>RR</i> is a 2-digit release number, and <i>M</i> is a 1-digit modification level. [Sequential = LST Non-Sequential = ID] [Switch = none]
svr_type	The type of database manager [Sequential = LST Non-Sequential = ID] [Switch = none]
svr_tz_disp	The number of seconds that the local time zone of the database manager is displaced from Greenwich mean time (GMT) [Sequential = LST, Non-Sequential = ID] [Switch = none]
svr_version	The version of the database manager [Sequential = LST Non-Sequential = ID] [Switch = none]
start_time	The timestamp of when the database manager was started [Sequential = LST Non-Sequential = ID] [Switch = none]
status	The status of the database manager. A value of Inactive will be stored if the agent cannot attach to the instance. Otherwise, a value of Active will be stored. [Sequential = LST Non-Sequential = ID] [Switch = none]

System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID] [Switch = none]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID] [Switch = none]

3.9. Table Statistics

Detailed information about the tables within each DB2 database is stored within the DB2.Table Detail table in the TeamQuest performance database. To obtain this information, the table DB2 monitor switch must be enabled within the DB2 Switches settings of the DB2 UDB Agent configuration entry for the instance. In addition, data for tables will only be collected for the databases listed in the Database List setting of the configuration entry. Records are only stored for those tables that were accessed within the interval.

Table Field Hierarchy

Class:	DB2
Subclass:	Table Detail
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /DB2/ <i>instancename</i>
TeamQuest Table Name:	DB2.Table Detail
Open Table Name:	DB2TABDETAIL
Collection interval:	N/A
Default retention:	1 month
Table type:	Event

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Non-Sequential = ID] [Switch = none]
database	The name of the database [Non-Sequential = ID] [Switch = table]
Instance	The name of the instance from which data is obtained. Up to 24 characters are displayed. [Non-Sequential = ID] [Switch = none]

Interval	The expected sampling interval in seconds [Non-Sequential = ID] [Switch = none]
overflows	The number of accesses (reads and writes) to overflowed rows of the table during the interval [Non-Sequential = SUM] [Switch = table]
reorgs	The number of page reorganizations executed for the table during the interval [Non-Sequential = SUM] [Switch = table]
rows_read	The number of rows read from the table during the interval. This value includes the number of accesses to overflowed rows. [Non-Sequential = SUM] [Switch = table]
rows_written	The number of rows written to the table during the interval [Non-Sequential = SUM] [Switch = table]
Sample_End_Time	The actual timestamp of when data collection for the sample completed [Non-Sequential = ID] [Switch = none]
schema	The schema of the table [Non-Sequential = ID] [Switch = table]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID] [Switch = none]
table_name	The name of the table [Non-Sequential = ID] [Switch = table]
table_type	The type of table, represented by one of the following one-character codes: C = Catalog D = Dropped R = Reorganization T = Temp [Non-Sequential = ID] [Switch = table]
Time	The timestamp of the data sample [Non-Sequential = ID] [Switch = none]

3.10. Tablespace Statistics

Detailed information about the tablespaces within each DB2 database is stored within the DB2.Tablespace Detail table in the TeamQuest performance database. To obtain this information, the bufferpool DB2 monitor switch must be enabled within the DB2 UDB Agent configuration entry for the instance. In addition, tablespace information will only be collected for the databases listed in the Database List setting of the configuration entry.

Table Field Hierarchy

Class:	DB2
Subclass:	Tablespace Detail
IT Resource Name:	/TeamQuest/System/systemname/DB2/instancename
TeamQuest Table Name:	DB2.Tablespace Detail
Open Table Name:	DB2TABSPACEDETAIL
Collection interval:	1 minute (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
a_data_rd_reqs	The number of asynchronous data read requests from the tablespace during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_data_rds	The number of data pages read asynchronously from the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_data_wrts	The number of data pages written asynchronously to the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_idx_rd_reqs	The number of asynchronous index read requests from the tablespace during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_idx_rds	The number of index pages read asynchronously from the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_idx_wrts	The number of index pages written asynchronously to the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]

a_rds	The number of data and index pages read asynchronously from the tablespace during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_rdtype	The amount of time in milliseconds spent reading pages asynchronously from the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_wrts	The number of data and index pages written asynchronously to the tablespace during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
a_wrttime	The amount of time in milliseconds spent writing pages asynchronously to the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID] [Switch = none]
avg_a_data_rds	The average number of data pages read per asynchronous data read request during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_a_idx_rds	The average number of index pages read per asynchronous index read request during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_a_rdtype	The average time in milliseconds spent reading a page asynchronously from the tablespace during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_a_wrttime	The average time in milliseconds spent writing a page asynchronously to the tablespace during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_rds	The average number of sectors that are read per direct read request during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]

avg_d_rdtype	The average time in milliseconds spent direct reading a sector from the tablespace during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_wrts	The average number of sectors that are written per direct write request during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_d_wrttime	The average time in milliseconds spent direct writing a sector to the tablespace during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_rdtype	The average time in milliseconds spent reading a page from the tablespace during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
avg_wrttime	The average time in milliseconds spent writing a page to the tablespace during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
content	The type of content in the tablespace, represented by one of the following one-character codes: A = Any L = Long S = System Temporary U = User Temporary [Sequential = ID Non-Sequential = ID] [Switch = none]
curr_pool_id	The identifier for the bufferpool currently being used by the tablespace [Sequential = LST Non-Sequential = ID] [Switch = none]
d_rd_reqs	The number of requests to perform direct read operations during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_rds	The number of read operations that do not use a bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_rdtype	The amount of time in milliseconds spent direct reading a sector from the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]

d_wrt_reqs	The number of requests to perform direct write operations during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_wrts	The number of write operations that do not use a bufferpool during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
d_wrttime	The amount of time in milliseconds spent direct writing a sector to the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
data_wrts	The number of data pages written to the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
database	The name of the database [Sequential = ID Non-Sequential = ID] [Switch = none]
extent_size	The extent size in pages used by the tablespace [Sequential = ID Non-Sequential = ID] [Switch = none]
files_closed	The number of times a file had to be closed to remain under the maxfilop configuration parameter when opening a new file for use from the tablespace during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
free_pgs	The number of pages currently free in the tablespace. This value is always <N/A> for system-managed space (SMS) tablespaces. [Sequential = LST Non-Sequential = SUM] [Switch = none]
hit_ratio	The percentage of logical reads that were satisfied for the tablespace without having to invoke a read from disk during the interval [Sequential = AVG Non-Sequential = AVG] [Switch = bufferpool]
idx_wrts	The number of index pages written to the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
Instance	The name of the instance from which data is obtained. Up to 24 characters are displayed. [Sequential = ID Non-Sequential = ID] [Switch = none]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID] [Switch = none]

<code>l_data_rds</code>	The number of logical read requests of data pages from the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
<code>l_idx_rds</code>	The number of logical read requests of index pages from the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
<code>l_rds</code>	The number of logical read requests of data and index pages from the tablespace during the interval [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
<code>min_recovery</code>	The timestamp of the earliest point in time to which a tablespace can be rolled forward. This value is not stored by default. [Sequential = LST Non-Sequential = ID] [Switch = none]
<code>next_pool_id</code>	The identifier for the bufferpool that will be used at the next database startup [Sequential = LST Non-Sequential = ID] [Switch = none]
<code>num_containers</code>	The number of containers in the tablespace. This value is not stored by default. [Sequential = LST Non-Sequential = SUM] [Switch = none]
<code>num_quiescers</code>	The number of users quiescing the tablespace. This value is not stored by default. [Sequential = LST Non-Sequential = SUM] [Switch = none]
<code>p_data_rds</code>	The number of physical read requests of data pages from the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
<code>p_idx_rds</code>	The number of physical read requests of index pages from the tablespace during the interval. Not stored by default. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
<code>p_rds</code>	The number of physical read requests of data and index pages from the tablespace during the interval. This includes reads done asynchronously. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]
<code>p_wrts</code>	The number of data and index pages written to the tablespace during the interval. This includes writes done asynchronously. [Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]

page_hwm	<p>The page in the tablespace that is holding the current high water mark. Essentially, this represents the page number of the first free extent following the last allocated extent of the tablespace. This value is always <N/A> for system-managed space (SMS) tablespaces.</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = none]</p>
page_size	<p>The page size in bytes used by the tablespace</p> <p>[Sequential = ID Non-Sequential = ID] [Switch = none]</p>
pend_free_pgs	<p>The number of pages in the tablespace that would become free if all pending transactions are committed or rolled back and new space is requested for an object. This value is always <N/A> for system-managed space (SMS) tablespaces.</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = none]</p>
prefetch_size	<p>The maximum number of pages the prefetcher retrieved from disk at a time for this tablespace</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
rdtime	<p>The amount of time in milliseconds spent reading pages from the tablespace during the interval. This value is not stored by default.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]</p>
rebal_last_moved	<p>The last extent moved by the rebalancer. This value is always <N/A> for system-managed space (SMS) tablespaces. This value is not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
rebal_mode	<p>The current state of rebalancing for the tablespace, represented by one of the following one-character codes:</p> <p style="margin-left: 40px;">N = No rebalancing F = Forward R = Reverse U = User Temporary</p> <p>This value is always <N/A> for system-managed space (SMS) tablespaces.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>
rebal_priority	<p>The priority at which the rebalancer is running in the database. This value is always <N/A> for system-managed space (SMS) tablespaces. This value is not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID] [Switch = none]</p>

rebal_processed	<p>The number of extents already moved by the rebalancer since it started or restarted. This value is always <N/A> for system-managed space (SMS) tablespaces. This value is not stored by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p>
rebal_remaining	<p>The number of extents to be moved by the rebalancer. This value is always <N/A> for system-managed space (SMS) tablespaces. This value is not stored by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = none]</p>
rebal_restart	<p>The timestamp of when a rebalancer was restarted after being paused or stopped. This value is always <N/A> for system-managed space (SMS) tablespaces. This value is not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID]</p> <p>[Switch = none]</p>
rebal_start	<p>The timestamp of when a rebalancer was initially started. This value is always <N/A> for system-managed space (SMS) tablespaces. This value is not stored by default.</p> <p>[Sequential = LST Non-Sequential = ID]</p> <p>[Switch = none]</p>
Sample_End_Time	<p>The actual timestamp of when data collection for the sample completed</p> <p>[Sequential = LST Non-Sequential = ID]</p> <p>[Switch = none]</p>
System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p> <p>[Switch = none]</p>
tablespace_name	<p>The name of the tablespace</p> <p>[Sequential = ID Non-Sequential = ID]</p> <p>[Switch = none]</p>
tablespace_state	<p>The current state of the tablespace</p> <p>[Sequential = LST Non-Sequential = ID]</p> <p>[Switch = none]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p> <p>[Switch = none]</p>
total_pgs	<p>The total number of pages in the tablespace. The bufferpool switch is not needed to collect this information on DMS tablespaces.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>[Switch = bufferpool]</p>
ts_id	<p>The unique identifier for the tablespace within the database</p> <p>[Sequential = ID Non-Sequential = ID]</p> <p>[Switch = none]</p>

ts_type	<p>The type of tablespace, represented by one of the following one-character codes:</p> <p>S = System Managed D = Database Managed</p> <p>[Sequential = ID Non-Sequential = ID] [Switch = none]</p>
unrd_pre_pgs	<p>The number of prefetched pages that were unread during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]</p>
usable_pgs	<p>The total number of pages in the tablespace minus overhead pages. The bufferpool switch is not needed to collect this information on DMS tablespaces.</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = bufferpool]</p>
used_pgs	<p>The number of pages currently used in the tablespace. The bufferpool switch is not needed to collect this information on DMS tablespaces.</p> <p>[Sequential = LST Non-Sequential = SUM] [Switch = bufferpool]</p>
wrttime	<p>The amount of time in milliseconds spent writing pages to the tablespace during the interval. Not stored by default.</p> <p>[Sequential = SUM Non-Sequential = SUM] [Switch = bufferpool]</p>

Section 4

Hewlett-Packard HP-UX Systems

Statistics for Hewlett-Packard HP-UX systems are collected by the TeamQuest collection agents.

This section contains a listing of the statistics collected for the system:

- System Activity Statistics (see 4.1)
- Disk Space Statistics (see 4.2)
- Network Statistics (see 4.3)
- Workload Statistics (see 4.4)
- Process Statistics (see 4.5)
- Hardware Inventory Statistics (see 4.6)
- System Log Statistics (see 4.7)
- General Log Statistics (see 4.8)
- TeamQuest Log Statistics (see 4.9)
- Derived Statistics (see 4.10)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

4.1. System Activity Statistics

The System Activity Agent is used to collect a wide variety of important system statistics. Major resources monitored by this agent include processors, memory, disks, network interfaces, and the operating system kernel.

Special Processing When Using Sequential Consolidation Method

Special processing occurs when certain records in the Block Device.by Device table are consolidated using the Sequential consolidation method. The following formulas are used to calculate the %busy, Actual_Interval, and record_count statistic values:

%busy

The %busy field uses a new consolidation method that uses the following formula to produce the consolidated %busy value:

$$\text{consolidated \%busy} = \%busy * \text{record_count} * \text{Actual_Interval}$$

At the end of the aggregation processing step after multiple records have been combined together to produce a single consolidated record, the %busy field contains the consolidated %busy value.

An additional processing step is performed using the following formula to produce a final %busy value that is stored into the consolidated record:

$$\%busy = \frac{\text{consolidated \%busy}}{\text{record_count} * \text{Interval}}$$

Note: The record_count field value used in the above formula must have already been generated using the record_count formula.

Block Device.by Device table records that have been stored by previous levels of TeamQuest collection agents do not contain the record_count field. For these records, a value of 1 is assumed for the record_count value.

Actual_Interval

For consolidated records (both reduced and not reduced), the Actual_Interval field should contain the Interval value at the end of the aggregation processing step.

record_count

The record_count field value is updated at the end of the aggregation processing step using the following formula:

$$\text{record_count} = \frac{(\sum(\text{record_count} * \text{Actual_Interval})) + (\text{Interval} - \sum \text{Actual_Interval})}{\text{Interval}}$$

The first part of the numerator is a summation of the record_count value multiplied by the Actual_Interval value across all of the consolidated records. This value should exist in the record_count field at the end of the aggregation processing step since the record_count field is weighted by Actual_Interval.

The Actual_Interval value used in the second part of the numerator is a summation of all the Actual_Interval values of the consolidated records. This value should exist in the Actual_Interval field at the end of the aggregation processing step since Actual_Interval is a summation.

The Interval value in the formula is the time range in seconds that the final consolidated record represents. For example, if 5-minute records are generated, the Interval value is calculated as 5 multiplied by 60.

Parameter Hierarchy

Class:	Block Device
Subclass:	by Device
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Block Device.by Device
Open Table Name:	BLKDEVBYDEVICE
Resource:	disk0, disk1, ...
Statistic Name:	
%busy	The percentage of time this device was servicing a transfer request [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-util.rpt
actq_avwait*	The average run queue wait time in milliseconds
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgresp*	The average response time of an I/O on a device. Calculated as await + avserv
avque	The average number of requests outstanding [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/io/dsk-q.rpt
avserv	The average time in milliseconds to service each transfer request (includes seek, rotation latency, and data transfer times) for the device [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-time.rpt
await	The average time in milliseconds that transfer requests are idle in the queue while the queue is occupied [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-time.rpt

Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
IO_intensity*	The activity of an I/O device. This is the product of the I/O response time in milliseconds and the I/O transfer rate in I/Os per second. This is proportional to the average queue length (the number of I/O requests waiting or in progress at the I/O device).
Kbytes/s	The rate at which data is transferred in kilobytes per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/hp-ux/sys-act/io/dsk-xfer.rpt /report/hp-ux/sys-act/io/top-dsk.rpt
record_count	The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record. [Sequential = AVG Non-Sequential = SUM]
reduction_name	The name of reduction rule [Sequential = ID Non-Sequential = ID]
reduction_source	The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H. [Sequential = ID Non-Sequential = ID]
transfers/s	The number of physical transfers to and from the disk per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/io/dsk-xfer.rpt
waitq_avwait*	The average wait queue wait time in milliseconds

Class:	Block Device
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Block Device.Summary
Open Table Name:	BLKDEVSUM
Statistic Name:	
transfers/s	The number of physical transfers to and from the disk per second [Sequential = AVG Non-Sequential = SUM]
Class:	CPU
Subclass:	by Processor
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.by Processor
Open Table Name:	CPUBYPROC
Resource:	cpu0, cpu1, ...
Statistic Name:	
%idle	The percentage of CPU time spent idle while no processes are waiting for I/O completion for this CPU [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/per-cpu.rpt
%nice	The percentage of CPU time spent running low priority user processes for this CPU [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/per-cpu.rpt
%sys	The percentage of CPU time spent running in system mode for this CPU [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/per-cpu.rpt
%usr	The percentage of CPU time spent running in user mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/per-cpu.rpt
%wio	The percentage of time spent idle while some processes are waiting for I/O completion [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/per-cpu.rpt
per_cpu_ticks	CPU clock ticks elapsed between two data samples [Sequential = SUM Non-Sequential = SUM]

Table Field Hierarchy

Class:	CPU
Subclass:	RelativePerformance
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.RelativePerformance
Open Table Name:	CPURELPERF
Collection interval:	1 minute
Default retentions:	1 month
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name

Description

Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
rel_unused	The amount of CPU resources not used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
rel_used	The amount of CPU resources used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Class:	CPU
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.Summary
Open Table Name:	CPUSUM
Statistic Name:	
%busy	The percentage of time the CPU was not idle [Sequential = AVG Non-Sequential = AVG]
%idle	The percentage of total CPU time spent idle while no processes are waiting for I/O completion [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/cpu-util.rpt
%nice	The percentage of total CPU time spent running low priority user processes [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/cpu-util.rpt
%sys	The percentage of total CPU time spent running in system mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/cpu-util.rpt
%usr	The percentage of total CPU time spent running in user mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/cpu-util.rpt
%wio	The percentage of time spent idle while some processes are waiting for I/O completion [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/cpu/cpu-util.rpt
online_cpus	The number of CPUs that were online at the end of the sampling interval [Sequential = LST Non-Sequential = SUM]

Class:	Kernel
Subclass:	Buffers
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Buffers
Open Table Name:	KNLBUFFS
Statistic Name:	
%rcache	The percentage of logical reads satisfied from the buffer cache [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/kernel/bufc-hit.rpt
%wcache	The percentage of logical writes satisfied from the buffer cache [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/kernel/bufc-hit.rpt
bread/s	The number of reads per second from devices into the buffer cache [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/bufc-xfr.rpt
bwrit/s	The number of writes per second from the buffer cache to devices [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/bufc-xfr.rpt
lread KB/s	The number of kilobytes (KB) read by processes from the system [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/bufc-kb.rpt
lread/s	The number of reads per second from the buffer cache to a process [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/bufc-xfr.rpt
lwrit/s	The number of write accesses of system buffers (logical writes) per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/bufc-xfr.rpt
lwrite KB/s	The number of kilobytes (KB) written to system buffers by processes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/bufc-kb.rpt
pread/s	The number of physical read requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/phys-xfr.rpt

pwr/s	The number of physical write requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/phys-xfr.rpt
readahead hits/s	The number of look ahead read requests that are satisfied by system buffers per second [Sequential = AVG Non-Sequential = SUM]
readahead KB/s	The total kilobytes (KB) read on lookahead read requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/bufc-kb.rpt
readahead/s	The number of reads from block devices to the kernel buffers per second, due to a lookahead request [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/bufc-xfr.rpt
Class:	Kernel
Subclass:	File Access
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.File Access
Open Table Name:	KNLFILEACCESS
Statistic Name:	
dirblk/s	The number of directory block reads issues per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/f-access.rpt
igets/s	The number of files located by i-node entry per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/f-access.rpt
namei/s	The number of file system path searches per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/f-access.rpt

Class: Kernel
Subclass: Interrupts
IT Resource Name: /TeamQuest/System/*systemname*/Kernel
TeamQuest Table Name: Kernel.Interrupts
Open Table Name: KNLITRPS
Statistic Name:
 device intr/s The number of hardware interrupts per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/sys-act/kernel/sys-intr.rpt
 trap/s The number of traps per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/sys-act/kernel/sys-intr.rpt

Class: Kernel
Subclass: IPC (interprocess communication)
IT Resource Name: /TeamQuest/System/*systemname*/Kernel
TeamQuest Table Name: Kernel.IPC
Open Table Name: KNLIPC
Statistic Name:
 msg/s The number of message operations (sends and receives) per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/sys-act/kernel/msg-sema.rpt
 sema/s The number of semaphore operations per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/sys-act/kernel/msg-sema.rpt

Class:	Kernel
Subclass:	Paging
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Kernel.Paging
Open Table Name:	KNLPAGING
Statistic Name:	
atch/s	<p>The number of page faults per second that are satisfied by reclaiming a page currently in memory (attaches per second)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/paging.rpt</p>
deficit	<p>An estimate of the pages needed by new swapped-in processes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/paging.rpt</p>
exfod/s	<p>The number of pages filled on demand from executables per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/pgfaults.rpt</p>
intransit_fault/s	<p>The number of intransit blocking page faults per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/pgfaults.rpt</p>
KB paged in/s	<p>The rate at which pages are paged in (expressed in kilobytes per second)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/page-amt.rpt</p>
KB paged out/s	<p>The rate at which pages are paged out (expressed in kilobytes per second)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/page-amt.rpt</p>
numexfod/s	<p>The number of times pages are filled on demand from executables per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/pgfaults.rpt</p>
numzfod/s	<p>The number of times pages are zero-filled on demand per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/pgfaults.rpt</p>
pg_fault/s	<p>The total number of page faults per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/pgfaults.rpt</p>

pg/fork	The average number of pages brought in as a result of each fork [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/syscall/fork-siz.rpt
pgdfree/s	The number of pages freed by the page daemon per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/pgfree.rpt
pgfrec/s	The number of pages reclaimed from the free list per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/pgreclms.rpt
pgfree/s	The total number of memory pages available on the free list per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/pgfree.rpt
pgin/s	The number of page-in requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/paging.rpt
pgout/s	The number of page-out requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/paging.rpt
pgrec/s	The total number of page reclaims per second (including pageout) [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/pgreclms.rpt
pgscan/s	The rate per second at which the page daemon scans pages to see if they can be freed [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/pgfree.rpt
rev/s	The number of revolutions per second of the paging daemon [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/pgfree.rpt
zfod/s	The number of pages zero-filled on demand per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/pgfaults.rpt

Class:	Kernel
Subclass:	Processes
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Processes
Open Table Name:	KNLPROCES
Statistic Name:	
active	<p>The count of active processes. This count is taken at the end of the sampling interval.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/process.rpt</p>
runnable	<p>The length of the run queue at the end of the sampling interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/process.rpt</p>
swapped	<p>A count of the runnable processes that have been swapped out. This count is taken at the end of the sampling interval.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/process.rpt</p>
Class:	Kernel
Subclass:	Queues
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Queues
Open Table Name:	KNLQS
Statistic Name:	
%runocc	<p>The percentage of time the run queue is occupied</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/q-util.rpt</p>
%swpocc	<p>The percentage of time the swap queue is occupied</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/q-util.rpt</p>
avg_cpuq_sz	<p>The average length of the run queue per CPU (a queue of processes in memory and runnable)</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
avg_runq_sz	<p>The average length of the run queue (a queue of processes in memory and runnable)</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
pswch/s	<p>The number of process switches per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/kernel/sys-act.rpt</p>

runq_sz	The average length of the run queue (a queue of processes in memory and runnable) while the run queue is occupied [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/kernel/q-sizes.rpt
swpq_sz	The average length of the swap queue (a queue of processes swapped out and ready to run) while the swap queue is occupied [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/q-sizes.rpt
Class:	Kernel
Subclass:	Swapping
IT Resource Name:	/TeamQuest/System/systemname/Memory
TeamQuest Table Name:	Kernel.Swapping
Open Table Name:	KNLSWAPPING
Statistic Name:	
freeswap	The average number of megabytes free for swapping from all configured swap areas. This includes swap space in the form of physical memory, but does not include any from file swap. [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/kernel/swapspc.rpt
KB swapped in/s	The rate at which pages are swapped in (expressed in kilobytes per second) [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/swap-amt.rpt
KB swapped out/s	The rate at which pages are swapped out (expressed in kilobytes per second) [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/swap-amt.rpt
swpin/s	The number of transfers to memory from swap per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/swaprate.rpt
swpot/s	The number of transfers from memory to swap per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/swaprate.rpt
totalswap	The total number of megabytes available for swapping from all configured swap areas. This includes swap space in the form of physical memory, but does not include any from file swap. [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/kernel/swapspc.rpt

Class:	Kernel
Subclass:	TTY
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.TTY
Open Table Name:	KNLTTY
Statistic Name:	
canch/s	The number of input characters processed by canon (canonical queue) per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/tty-xfer.rpt
inch/s	The rate at which characters are read from teletypewriter (TTY) devices per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/tty-xfer.rpt
mdmin/s	The number of modem interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/tty-intr.rpt
outch/s	The number of output characters transferred per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/tty-xfer.rpt
rcvin/s	The number of receiver hardware interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/tty-intr.rpt
xmtin/s	The number of transmitter hardware interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/tty-intr.rpt

Class:	Load Average
Subclass:	by Processor
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	Load Average.by Processor
Open Table Name:	LOADAVGBYPROC
Resource:	cpu0, cpu1, ...
Statistic Name:	
1 min	The number of processes in the run queue averaged over the last 1 minute. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/load-by.rpt
5 min	The number of processes in the run queue averaged over the last 5 minutes. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/load-by.rpt
15 min	The number of processes in the run queue averaged over the last 15 minutes. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/load-by.rpt
Class:	Load Average
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Load Average.Summary
Open Table Name:	LOADAVGSUM
Statistic Name:	
1 min	The number of processes in the run queue averaged over the last 1 minute. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/load-avg.rpt
5 min	The number of processes in the run queue averaged over the last 5 minutes. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/load-avg.rpt
15 min	The number of processes in the run queue averaged over the last 15 minutes. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/kernel/load-avg.rpt

Class:	Memory
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Memory
Open Table Name:	MEM
Statistic Name:	
%realinuse	<p>The actual usage of real memory as a percentage of total real memory in use</p> <p>[Sequential = LST Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/memory/musage.rpt</p>
%virtualinuse	<p>The actual usage of virtual memory as a percentage of total virtual memory in use</p> <p>[Sequential = LST Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/memory/musage.rpt</p>
bufcache	<p>The amount of memory being used by the buffer cache in kilobytes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/memory/bufcach.rpt</p>
freemem	<p>The amount of free memory in kilobytes</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/hp-ux/sys-act/memory/freemem.rpt</p> <p>/report/hp-ux/sys-act/memory/memory.rpt</p>
physmem	<p>The total amount of physical memory in kilobytes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/memory/memory.rpt</p>
real	<p>The total amount of real (or physical) memory in use in kilobytes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/memory/memory.rpt</p>
virtual	<p>The total amount of virtual memory in use in kilobytes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/memory/memory.rpt</p>

Class:	System Call
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	System Call.Summary
Open Table Name:	SYSCALLSUM
Statistic Name:	
exec/s	The number of exec system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/syscall/imp-scal.rpt
fork/s	The number of fork system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/syscall/imp-scal.rpt
rchar/s	The number of characters transferred by read system calls in the interval in bytes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/syscall/scal-xfr.rpt
scall/s	The number of system calls made per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/syscall/imp-scal.rpt
sread/s	The number of read system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/syscall/imp-scal.rpt
swrit/s	The number of write system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/syscall/imp-scal.rpt
wchar/s	The number of characters transferred by write system calls in the interval in bytes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/syscall/scal-xfer.rpt

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
bsp interval	The number of seconds elapsed between two data samples of the System Activity Agent [Sequential = SUM Non-Sequential = ID]
elapsed_ticks	The number of CPU clock ticks elapsed between two data samples [Sequential = SUM Non-Sequential = SUM]
tqbsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqbsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name

Description

Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

4.2. Disk Space Statistics

The Disk Space Agent tracks the disk space usage of locally mounted file systems. However, any data about Network File System (NFS) mounted file systems are not collected by the agent because the loss of the network connection to the NFS may result in hung processes.

Note: *If the file system names are longer than 52 characters, problems may occur.*

Parameter Hierarchy

Class:	Disk Space
Subclass:	by File System
IT Resource Name:	/TeamQuest/System/systemname/Disk
TeamQuest Table Name:	Disk Space.by File System
Open Table Name:	DISKSPACEBYFILESYS
Resource:	file system1, file system2, ...
Statistic Name:	
%inodes free*	The percentage of i-nodes available (not in use) on the file system at the end of the interval View Report: /report/hp-ux/dskspace/total/pct-inod.rpt
%inodes used*	The percentage of i-nodes in use on the file system at the end of the interval View Report: /report/hp-ux/dskspace/total/pct-inod.rpt
%space free*	The percentage of total space available (not in use) on the file system at the end of the interval View Reports: /report/hp-ux/dskspace/total/low-ones.rpt /report/hp-ux/dskspace/total/pctspace.rpt
%space used*	The percentage of total space in use on the file system at the end of the interval View Reports: /report/hp-ux/dskspace/total/fullest.rpt /report/hp-ux/dskspace/total/pctspace.rpt
%user space free*	The percentage of total user space available (not in use) on the file system at the end of the interval View Reports: /report/hp-ux/dskspace/user/low-ones.rpt /report/hp-ux/dskspace/user/pctspace.rpt
%user space used*	The percentage of total user space in use on the file system at the end of the interval View Reports: /report/hp-ux/dskspace/user/fullest.rpt /report/hp-ux/dskspace/user/pctspace.rpt

capacity	The percentage of total space in use on the file system at the end of the interval [Sequential = LST Non-Sequential = AVG]
free (Mb)	The amount of space available (not in use) on the file system in megabytes at the end of the interval including the space held back from normal users [Sequential = LST Non-Sequential = SUM] View Reports: /report/hp-ux/dskspace/total/dskspace.rpt /report/hp-ux/dskspace/total/low-ones.rpt
free inodes	The number of available (not in use) i-nodes on the file system at the end of the interval [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/dskspace/total/i-nodes.rpt
total (Mb)	The total (used + available) amount of space on the file system in megabytes at the end of the interval including the space held back from normal users [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/dskspace/total/dskspace.rpt
total inodes	The total (used + available) number of i-nodes on the file system at the end of the interval [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/dskspace/total/i-nodes.rpt
user free (Mb)	The amount of space available (not in use) on the file system in megabytes (Mb) at the end of the interval not including the space held back from normal users [Sequential = LST Non-Sequential = SUM] View Reports: /report/hp-ux/dskspace/user/dskspace.rpt /report/hp-ux/dskspace/user/low-ones.rpt
user total (Mb)*	The total (used + available) amount of space on the file system in megabytes (Mb) at the end of the interval not including the space held back from normal users View Report: /report/hp-ux/dskspace/user/dskspace.rpt

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
dsp interval	The number of seconds elapsed between two data samples of the Disk Space Agent [Sequential = SUM Non-Sequential = ID]
tqdsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqdsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name

Description

Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

4.3. Network Statistics

The Network Agent collects data about the interaction of the system with the network. The types of data collected by the agent include network interfaces, NFS and Remote Protocol Call (RPC), and Transmission Control Protocol (TCP).

In NFS environments, the term *server* refers to the system that owns a file system and allows other systems on the network to access the files on the file system. A *client* is a system that mounts a nonlocal file system such as NFS and accesses files from the NFS mounted file system. Any given system could be a server to some file systems and a client to other file systems. It is not uncommon to have systems that export several file systems for other systems to use and mount several other network file systems. All the NFS statistics collected by this agent show the activity of NFS clients and servers that exist on the system where the agent is run.

Parameter Hierarchy

Class:	Network
Subclass:	by Interface
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	Network.by Interface
Open Table Name:	NETBYINTERFACE
Resource:	interface0, interface1, ...
Statistic Name:	
collisions/s	The number of network collisions per second on Carrier Sense Multiple Access (CSMA) interfaces [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/net-errs.rpt
in errors/s	The number of network input errors per second. The value is always 0 on systems running HP-UX level 11.00. [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/net-errs.rpt
in packets/s	The number of network input packets per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/net-pkts.rpt
out errors/s	The number of network output errors per second. The value is always 0 on systems running HP-UX level 11.00. [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/net-errs.rpt
out packets/s	The number of network output packets per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/net-pkts.rpt

Class:	Network
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	Network.Summary
Open Table Name:	NETSUM
Statistic Name:	
errors/s	The total number of network errors per second for all network interfaces on the system. The value is always 0 on systems running HP-UX level 11.00. [Sequential = AVG Non-Sequential = SUM]
in packets/s	The total number of network input packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/net-sum.rpt
out packets/s	The total number of network output packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/net-sum.rpt
total packets/s	The total number of network packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]
Class:	NFS
Subclass:	Client
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFS.Client
Open Table Name:	NFSCLI
Statistic Name:	
badcalls/s	The total number of Network File System (NFS) calls per second rejected by the client [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/client.rpt
calls/s	The total number of NFS calls sent by the client per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/client.rpt
gets/s	The total number of times per second an NFS client handle was received [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/netowrk/client.rpt

Class: NFS
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFS.Server
Open Table Name: NFSSERV
Statistic Name:
 badcalls/s The total number of NFS calls per second rejected by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/network/server.rpt
 calls/s The total number of NFS calls per second received by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/network/server.rpt

Class: NFSv2
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv2.Client
Open Table Name: NFSV2CLI
Statistic Name:
 calls/s* The number of NFS version 2 calls per second sent by the client
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/network/clnt-v2.rpt

Class:	NFSv2
Subclass:	Client
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFSv2.Client
Open Table Name:	NFSV2CLI
Resource:	create, getattr, link, lookup, mkdir, null, read, readdir, readlink, remove, rename, rmdir, root, setattr, statfs, symlink, write, writecache
Statistic Name:	
reqs/s	The number of NFS version 2 requests per second sent by the client [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/clnt-v2.rpt
Class:	NFSv2
Subclass:	Server
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFSv2.Server
Open Table Name:	NFSV2SERV
Statistic Name:	
calls/s*	The number of NFS version 2 calls per second received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/srvr-v2.rpt
Class:	NFSv2
Subclass:	Server
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFSv2.Server
Open Table Name:	NFSV2SERV
Resource:	create, getattr, link, lookup, mkdir, null, read, readdir, readlink, remove, rename, rmdir, root, setattr, statfs, symlink, write, writecache
Statistic Name:	
reqs/s	The number of NFS version 2 requests per second received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/network/srvr-v2.rpt

Class:	NFSv3
Subclass:	Client
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFSv3.Client
Open Table Name:	NFSV3CLI
Statistic Name:	
calls/s*	<p>The number of NFS version 3 calls per second sent by the client. NFS version 3 statistics are only available for systems running HP-UX level 11.0.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/clnt-v3.rpt</p>
Class:	NFSv3
Subclass:	Client
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFSv3.Client
Open Table Name:	NFSV3CLI
Resource:	access, commit, create, fsinfo, fsstat, getattr, link, lookup, mkdir, mknod, null, pathconf, read, readdir, readdir+, readlink, remove, rename, rmdir, setattr, symlink, write
Statistic Name:	
reqs/s	<p>The number of NFS version 3 requests per second sent by the client. NFS version 3 statistics are only available for systems running HP-UX level 11.0.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/clnt-v3.rpt</p>
Class:	NFSv3
Subclass:	Server
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFSv3.Server
Open Table Name:	NFSV3SERV
Resource:	access, commit, create, fsinfo, fsstat, getattr, link, lookup, mkdir, mknod, null, pathconf, read, readdir, readdir+, readlink, remove, rename, rmdir, setattr, symlink, write
Statistic Name:	
reqs/s	<p>The number of NFS version 3 requests per second received by the server. NFS version 3 statistics are only available for systems running HP-UX level 11.0.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/srvr-v3.rpt</p>

Class: NFSv3
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Server
Open Table Name: NFSV3SERV
Statistic Name:
 calls/s* The number of NFS version 3 calls per second received by the server.
 NFS version 3 statistics are only available for systems running HP-UX
 level 11.0.
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/network/srvr-v3.rpt

Class: RPC
Subclass: Client.Connectionless
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: RPC.Client.Connectionless
Open Table Name: RPCCLICONNLESS
Statistic Name:
 badcalls/s The number of connectionless RPC calls per second rejected by the
 server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/network/rpc/clnt-cl.rpt

 badxid/s The number of times per second a reply from a server was received that
 did not correspond to any outstanding connectionless RPC call
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/network/rpc/clnt-cl.rpt

 calls/s The total number of connectionless RPC calls per second sent by the
 client
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/network/rpc/clnt-cl.rpt

 newcred/s The number of times per second connectionless RPC authentication
 information had to be refreshed by the client
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/hp-ux/network/rpc/clnt-cl.rpt

retrans/s	<p>The number of times per second a connectionless RPC call had to be retransmitted by the client due to a timeout while waiting for a reply from the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/rpc/clnt-cl.rpt</p>
timeouts/s	<p>The number of times per second a connectionless RPC call timed out while waiting for a reply from the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/rpc/clnt-cl.rpt</p>
Class:	RPC
Subclass:	Server.Connectionless
IT Resource Name:	/TeamQuest/System/systemname/Network
TeamQuest Table Name:	RPC.Server.Connectionless
Open Table Name:	RPCSERVCONNLESS
Statistic Name:	
badcalls/s	<p>The number of connectionless RPC calls per second rejected by the server. The sum of badlens/s and xdrcalls/s.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/rpc/srvr-cl.rpt</p>
badlens/s	<p>The number of connectionless RPC calls per second received by the server with a length shorter than a minimum-sized RPC call</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/rpc/srvr-cl.rpt</p>
calls/s	<p>The number of connectionless RPC calls per second received by the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/rpc/srvr-cl.rpt</p>
nullrecvs/s	<p>The number of times per second a connectionless RPC call was not available when it was thought to be received by the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/rpc/srvr-cl.rpt</p>
xdrcalls/s	<p>The number of connectionless RPC calls per second by the server whose header could not be External Data Representation (XDR) decoded</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/network/rpc/srvr-cl.rpt</p>

Class:	TCP
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	TCP
Open Table Name:	TCP
Statistic Name:	
ftp conn	<p>The number of ftp connections at the end of the interval. Not available for systems running HP-UX level 11.0. [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/network/tcp/connect.rpt</p>
http conn	<p>The number of HTTP connections at the end of the interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/network/tcp/connect.rpt</p>
other conn	<p>The number of other connections at the end of the interval. Not available for systems running HP-UX level 11.0. [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/network/tcp/connect.rpt</p>
rlogin conn	<p>The number of rlogin connections at the end of the interval. Not available for systems running HP-UX level 11.0. [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/network/tcp/connect.rpt</p>
telnet conn	<p>The number of telnet connections at the end of the interval. Not available for systems running HP-UX level 11.0. [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/network/tcp/connect.rpt</p>
total conn	<p>The total number of connections at the end of the interval. Not available for systems running HP-UX level 11.0. [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/network/tcp/connect.rpt</p>

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
net interval	The number of seconds elapsed between two data samples of the Network Agent [Sequential = SUM Non-Sequential = ID]
tqbnp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqbnp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

4.4. Workload Statistics

Workload statistics are maintained in the TeamQuest performance database by the Process-Workload Agent. The statistics are classified by the hierarchy of key names.

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
twarp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
twarp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]
warp interval	The number of seconds elapsed between two data samples of the Process-Workload Agent [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Parameter Hierarchy

Class:	Workload
Subclass:	by Workload
IT Resource Name:	<i>/TeamQuest/System/systemname/workload</i> <i>/workloadset/workload</i>
TeamQuest Table Name:	Workload.by Workload
Open Table Name:	WLBYWORKLOAD
Workload Set:	WLS1, WLS2, ...
Workload:	WL1, WL2, ...
Statistic Name:	
%cpu	The percentage of total CPU consumed by the workload. Total CPU time is the value of the sampling interval multiplied by the number of CPUs on the system. Thus, if the sum of %cpu for all workloads is less than 100%, some CPUs must have been idle for some time during the sampling interval. [Sequential = AVG Non-Sequential = SUM] View Report: <i>/report/hp-ux/workload/pct-cpu.rpt</i>
avgmem	The cumulative swap process image size in kilobytes of all the running processes in the workload at the end of the sampling interval [Sequential = AVG Non-Sequential = SUM]
etime	The sum of the elapsed times in seconds of all the processes in the workload. Dividing this number by the number of processes in the workload (ponging + pcomplete) equals the average time a process in the workload existed during the sampling interval. [Sequential = SUM Non-Sequential = SUM]
lioch	The number of logical characters in kilobytes transferred by the workload during the sampling interval [Sequential = SUM Non-Sequential = SUM] View Report: <i>/report/hp-ux/workload/lioch.rpt</i>
majflt	The number of major page faults generated by the workload for processes that were active at the end of the sampling interval. A major page fault is a page fault that requires I/O. [Sequential = SUM Non-Sequential = SUM] View Report: <i>/report/hp-ux/workload/maj-flt.rpt</i>
pcomplete	The number of processes completed in the sampling interval. For process data, the same number is called cproc. [Sequential = SUM Non-Sequential = SUM] View Report: <i>/report/hp-ux/workload/num-proc.rpt</i>

pio	<p>The number of physical I/O transfers done by the workload during the sampling interval. The number reported represents only the completed processes during the sampling interval. [Sequential = SUM Non-Sequential = SUM] View Report: /report/hp-ux/workload/pio.rpt</p>
pongoing	<p>The number of processes running at the end of the sampling interval. In process data, the same number is called oproc. The number of processes in a workload could be derived by adding ponging and pcomplete. This sum is called nproc in process data. [Sequential = LST Non-Sequential = SUM] View Report: /report/hp-ux/workload/num-proc.rpt</p>
prss	<p>The resident set size in kilobytes of private memory occupied by all the running processes in the workload at the end of the sampling interval [Sequential = AVG Non-Sequential = SUM]</p>
pstart	<p>The number of processes started in the sampling interval. In process data, this number is called sproc. [Sequential = SUM Non-Sequential = SUM] View Report: /report/hp-ux/workload/num-proc.rpt</p>
record_count	<p>The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record. [Sequential = AVG Non-Sequential = SUM]</p>
reduction_name	<p>The name of reduction rule [Sequential = ID Non-Sequential = ID]</p>
reduction_source	<p>The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H. [Sequential = ID Non-Sequential = ID]</p>
rss	<p>The resident set size (real memory) in kilobytes of all processes running at the end of the sampling interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>. [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/workload/rss.rpt</p>
srss	<p>The resident set size in kilobytes of shared memory occupied by all the running processes in the workload at the end of the sampling interval [Sequential = AVG Non-Sequential = SUM]</p>

syscpu	<p>The system CPU time in seconds used by the workload. System CPU time is the time spent running in kernel mode (for example, the time spent in executing system calls, paging, and so on).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/workload/sys-cpu.rpt</p>
threads	<p>The number of threads at the end of the sampling interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/workload/threads.rpt</p>
totcpu	<p>The total CPU time in seconds used by the workload. This value is the same as the sum of usrcpu and syscpu.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/workload/cpu-util.rpt</p>
usrcpu	<p>The user CPU time in seconds used by the workload. User CPU time is the time spent running in user mode.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/workload/user-cpu.rpt</p>

4.5. Process Statistics

The Process-Workload Agent collects process data from the operating system and processes accounting files. The Process-Workload Agent calculates the usage of every process in a given interval, applies reduction definitions to each process, and stores the reduced process data. It also applies workload definitions to the reduced process data and stores system resource usage by workload.

Reduced Process Records

The Process-Workload Agent collects process data and reduces the data according to the user-defined reduction definitions. A reduction definition may cause multiple processes to be merged into a single record. Thus, a process record contains data about one or more processes. When you are looking at the resource usage numbers, it is important to know how many processes a process record actually represents. The `nproc` data item indicates exactly how many processes each process record is representing. When a process record is representing more than one process, the resource usage fields such as `totcpu`, `rss`, and `pio_t` are the sum of the resource usage of the individual processes. When all the processes do not have the same value for a field, the identifier fields such as `command`, `login`, and `gid` are set to `<Multi>`. When data for some fields is not available, the fields are set to `<N/A>`.

Disabling Reduction Definitions

If you want to look at the details of every individual process and do not wish to have merged process records, you must disable reduction processing by making all reduction sets inactive. However, with reduction processing disabled, more records have to be stored and more disk space is needed. For information on disabling reductions, see the *TeamQuest Performance Software Administration Guide*.

Process Data with Process Accounting Turned Off

If process accounting is turned off, the process data is incomplete as data about completed processes is not available. In this case, the process data only shows a portion of the activity in the sampling interval. To find out if process accounting is turned on or off, look at the cproc field of all the process records. If the cproc field indicates all zeros, it means that processes were not completed in that sampling interval or that the process accounting is turned off.

When process accounting is not turned on, the process record, <Other> includes the CPU time for processes that completed during the interval.

For more information on process accounting, see the *TeamQuest Performance Software Administration Guide*.

Retrieving Hardware Configuration Information

The Process-Workload Agent retrieves hardware configuration information. The information is stored upon start up and once-a-day in the HINV.Summary, HINV.Devices, and HINV.FileSystem table files of the TeamQuest performance database. The information is also stored if the agent detects a change in configuration.

Table Field Hierarchy

Class:	HPUX
Subclass:	Process
IT Resource Name:	/TeamQuest/System/systemname/Process
TeamQuest Table Name:	HPUX.Process
Open Table Name:	HPUXPROC
Collection interval:	Based on the primary aggregation set
Default retention:	1 day
Table type:	Performance

Note: The collection interval is also dependent on the Processes Only setting in the configuration file for the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the *TeamQuest Performance Software Administration Guide*.

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgmemb_t	The size of the swappable process image in kilobytes. If a process starts and ends within the same interval, the number is unavailable and is marked as <N/A>. [Sequential = AVG Non-Sequential = SUM]

btime	The start time of the process. For process records representing more than one process, this field shows the earliest of the start times. [Sequential = FST Non-Sequential = ID]
command	The command name of the process. If a process starts and ends within the same interval, only up to 8 characters of the command name can be displayed. Otherwise, up to 14 characters are displayed. Therefore, an “automountd” process may appear as “automoun” if it starts and ends within the same interval. In such cases, you may want to set up your workload, reduction, or filter definitions to catch both the command names as in the following: command = {“automoun” , “automountd”} [Sequential = ID Non-Sequential = ID]
cproc	The number of processes completed in the interval [Sequential = SUM Non-Sequential = SUM]
etime	The elapsed time in seconds for the current interval. This number tells how long a process existed in the current interval. [Sequential = SUM Non-Sequential = SUM]
etime_t	The total elapsed time in seconds. This number tells how long a process existed since it started. For a single process, this is the same as the value of the Time field minus the value of the btime field. [Sequential = LST Non-Sequential = SUM]
fullcmd	<p>The full command string, including arguments, for the process. If a process begins and ends in the sampling interval, the field is unavailable and is marked as <N/A>. This is an important consideration when using a workload, reduction, or filter definition. The number of characters that are stored is determined by the Command Length setting of the Process-Workload Agent configuration file in TeamQuest Manager. You can also have either the first or the last <i>N</i> characters of the command displayed. The setting is specified by the Command Orientation setting of the Process-Workload Agent configuration file in TeamQuest Manager. The limit of characters available from the operating system data source is 1020.</p> <p>For a description of how this may affect collection when the LAST Command Orientation is configured for the Process-Workload Agent, see the <i>TeamQuest Performance Software Administration Guide</i>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
gid	The real group identifier of the owner of the process [Sequential = ID Non-Sequential = ID]
group	The group name of the owner of the process. This field is derived from the gid. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
lioch	The number of logical characters in kilobytes transferred in the current interval [Sequential = SUM Non-Sequential = SUM]

lioch_t	The total number of logical characters in kilobytes transferred since the process started [Sequential = LST Non-Sequential = SUM]
login	The login name of the owner of the process. This field is derived from the uid. [Sequential = ID Non-Sequential = ID]
majflt	The number of major page faults generated in the current interval. A major page fault is a page fault that requires I/O. If a process ends in the sampling interval, the number is unavailable and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
majflt_t	The total number of major page faults generated since the process started. A major page fault is a page fault that requires I/O. If a process ends in the sampling interval, the number is unavailable and is marked as <N/A>. [Sequential = LST Non-Sequential = SUM]
nproc	The number of processes that this process record represents. If a process starts and ends within the same interval, the number is unavailable and is marked as <N/A>. In a reduced record, it is the number of processes that were merged together to form a single process record. When no time consolidation is applied to the process data, this number should equal the sum of cproc and oproc. When consolidating over time, the number represents the average number of processes in the process record for the consolidated period. An average is used because processes can move in and out and between reduced records from interval to interval. [Sequential = AVG Non-Sequential = SUM]
oproc	The number of ongoing processes at the end of the interval [Sequential = LST Non-Sequential = SUM]
pctcpu	The percentage of total available CPU time the process used in the current sampling interval [Sequential = AVG Non-Sequential = SUM]
pid	The process identifier number. If a process starts and ends within an interval, this number is unavailable and is marked as <N/A>. [Sequential = ID Non-Sequential = ID]
pio	The number of physical I/O requests for the current interval. The number reported represents only the completed processes during the sampling interval. [Sequential = SUM Non-Sequential = SUM]
pio_t	The total number of physical I/O requests since the process started. The number reported represents only the completed processes during the sampling interval. [Sequential = LST Non-Sequential = SUM]
ppid	The numerical identifier of the parent of a process. If a process starts and ends within an interval, the number is unavailable and is marked as <N/A>. [Sequential = ID Non-Sequential = ID]

pri	<p>The priority of the process. Higher numbers mean lower priority. If a process ends within an interval, the number is unavailable and is marked as <N/A>. If the process record represents more than one process, the priority is an average of the constituent running processes' priority values.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
prmid	<p>The Process Resource Manager (PRM) identifier for the process. This value is only available after the PRM has been installed and configured on your system. Otherwise the value is <N/A>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
prss	<p>The resident set size in kilobytes of private memory for the process at the end of the sampling interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
redname	<p>The reduction name of the process record. If a process did not match any of the reduction definitions, then it would not be reduced and will not have a reduction name.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
rss	<p>The resident set size (real memory size) of the process at the end of the interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
sproc	<p>The number of processes started in the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
srss	<p>The resident set size in kilobytes of shared memory occupied by all the running processes at the end of the sampling interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
syscpu	<p>The system CPU time in seconds for the current interval. System CPU time is the time spent running in kernel mode (for example, the time spent in executing system calls, paging, and so on). If an application is taking a lot of syscpu time, you may want to optimize the use of system calls (for example, use a larger block size for I/O).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
syscpu_t	<p>The total system CPU time in seconds</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>

threads	<p>The number of threads associated with the process at the end of the interval. This value represents the number of lightweight processes (LWPs) or kernel-supported user threads. A thread is a dynamic object that represents a control point in a process and executes a sequence of instructions.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
totcpu	<p>The total CPU time in seconds used in the current interval. This number is the same as the sum of usrcpu and syscpu.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
totcpu_t	<p>The total CPU time (user + system) in seconds used by the process since it started. This number is the same as the sum of usrcpu_t and syscpu_t.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
tty	<p>The controlling terminal identifier in dev_t format. For the processes without a controlling terminal, this field will contain a -1.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
ttyname	<p>The controlling terminal for the process. It is a device name without the /dev/ prefix. This is derived from tty. For the processes without a controlling terminal, this field will contain a question mark (?).</p> <p>[Sequential = ID Non-Sequential = ID]</p>
uid	<p>The real user id of the process owner</p> <p>[Sequential = ID Non-Sequential = ID]</p>
usrcpu	<p>The user CPU time in seconds for the current interval. User CPU time is the time spent running in user mode. If an application is taking a lot of usrcpu time, you should try to optimize the code, if possible.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
usrcpu_t	<p>The total user CPU time in seconds since the start of the process</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Workload	<p>The workload set and the workload associated with the process. When the Process-Workload Agent stores the process record, this field is blank. When process records are reported, the workload can be evaluated and is shown in the report.</p> <p>This field is available for reporting only when using TeamQuest Analyzer and TeamQuest tView.</p> <p>Workload evaluation takes place when data is retrieved from the database, based on workload sets defined in the database where the data is stored. Workload sets reported in the Workload field do not have to be active.</p> <p>For more information on workload evaluation, see the <i>TeamQuest Analyzer User Guide</i> or <i>TeamQuest Performance Software Command Line Interfaces Reference Manual</i>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>

This field is available for reporting only when using TeamQuest View or TeamQuest cView.
[Sequential = ID Non-Sequential = ID]

The hardware inventory statistics listed in this section are stored in the TeamQuest performance database tables by the Process-Workload Agent.

Note: *The storage of hardware inventory records depends on the Hardware Inventory setting in the configuration file of the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.*

Class:	HINV
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.Summary
Open Table Name:	HINVSUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

TQ-40023.4

cpu_count	The number of configured processors [Non-Sequential = ID]
cpu_speed	The speed of the processor in MHz or GHz [Non-Sequential = ID]
cpu_type	The basic instruction set architecture of the current system [Non-Sequential = ID]
logical_cpu_count	The number of logical processors [Non-Sequential = ID]
mem_size	The size of configured random access memory in kilobytes, where 1 kilobyte = 1,024 bytes [Non-Sequential = ID]
memory	The size of configured random access memory in megabytes, where 1 megabyte = 1,048,576 bytes [Non-Sequential = ID]
memory_size	The size of configured random access memory in megabytes or gigabytes [Non-Sequential = ID]
model	The name of the hardware implementation or platform [Non-Sequential = ID]
os_release	The name and level of this implementation of the operating system [Non-Sequential = ID]
pagesize	The size of a page of memory [Non-Sequential = ID]
partition_type	The partition type of the system. The value indicates the system hypervisor type, guest type, logical partition type, zone type, or logical domain type. If the system does not have a partition type, this field will be blank. [Non-Sequential = ID]
serial	The hardware-specific serial number of the physical machine [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_identifier	The information used to identify the system [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
timezone	The time zone where the data was collected [Non-Sequential = ID]
TQLevel	The level of TeamQuest Manager [Non-Sequential = ID]

The HINV.CPUModel table stores best-match, relative performance data about the system configuration. This table is created by the Hardware Inventory Agent (**tqhin**) to map physical hardware to a CPU model that describes performance in relative terms. This table is not created for any virtualized system. It is populated for physical systems only. It is not populated for VMware guests, Hyper-V guests, Solaris logical domains (LDMs), Solaris guest LDMs, KVM guests, and Linux on POWER systems.

Table Field Hierarchy

Class:	HINV
Subclass:	CPUModel
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	HINV.CPUModel
Open Table Name:	HINVCPUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
cpu_chips	The number of CPU chips or sockets [Non-Sequential = ID]
cpu_confidence	The percentage of confidence in the correctness of the CPU match based on model, frequency, and configuration (chips, cores, threads) [Non-Sequential = SUM]
cpu_cores	The number of CPU cores or processors on an individual CPU chip [Non-Sequential = ID]
cpu_name	The name of the selected CPU [Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Non-Sequential = ID]
cpu_speed	The speed of the processor in megahertz (MHz) or gigahertz (GHz) [Non-Sequential = ID]
cpu_threads	The number of CPU threads on an individual CPU core or processor [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]

Time	The timestamp of the data sample [Non-Sequential = LST]
user_override	The user override status of the default TeamQuest generated CPU match. This field is not currently used and should appear as 0. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	CPU Thread Speeds
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.CPU Thread Speeds
Open Table Name:	HINVCPU THREADSPEEDS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
speed_up_factor	The performance improvement when there are multiple active threads per core, compared to when there is only one active thread per core [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
thread_number	The number of active threads [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	Devices
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.Devices
Open Table Name:	HINVDEVS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
class	The device classification: controller, disk, or tape [Non-Sequential = ID]
controller	The device path indicator which defines a connection to another device [Non-Sequential = ID]
lun_id	The globally unique Logical Unit Number (LUN) identifier for Storage Area Network (SAN) based disk devices. This field is blank for non-SAN based disk devices, CD-ROM drives, tape drives, and so on. [Non-Sequential = ID]
name	The unique identifier for this device [Non-Sequential = ID]
name2	The alternate device name. This field may be blank. [Non-Sequential = ID]
product	The product identifier. This field may be blank. [Non-Sequential = ID]
revision	The revision level for this product. This field may be blank. [Non-Sequential = ID]
rpm	The speed at which the media spins. If an actual value cannot be obtained for the device, a default value of 7,200 is used. [Non-Sequential = ID]
swap	A true or false statement which indicates whether or not a swap file exists on the device [Non-Sequential = ID]
sequence	The sequence number of the device [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]

Time	The timestamp of the data sample [Non-Sequential = ID]
vendor	The name of the device vendor. This field may be blank. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	FileSystem
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.FileSystem
Open Table Name:	HINVFILESYS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
BlkSize	The size of a block on the file system [Non-Sequential = ID]
Device	The path for the device on which the file system is mounted [Non-Sequential = ID]
Name	The unique identifier for the file system [Non-Sequential = ID]
Source	The source physical disk or logical volume of the file system. This field is always blank for this platform. [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
TotBlks	The total number of blocks on the file system [Non-Sequential = ID]
TotFiles	The maximum total number of files, as represented by inodes, possible on the file system. Some inodes may be used for entities other than visible files. [Non-Sequential = ID]
TotSize	The total amount of space on the file system in megabytes [Non-Sequential = ID]
Type	The type of the file system [Non-Sequential = ID]

4.7. System Log Statistics

The System Log Agent is used to collect system log messages generated by the system log daemon (syslogd). The System Log Agent stores these messages in the TeamQuest performance database for analysis and alarm reporting. The log messages are separated into four fields; the time that the message was posted, the host system from where the message was initiated, the program or user that posted the message, and the text of the message.

Table Field Hierarchy

Class:	System
Subclass:	System Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /System Log
TeamQuest Table Name:	System.System Log
Open Table Name:	SYSSYSTEMLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Event_Time	The time that the message was logged to the system log [Non-Sequential = ID]
Loghost	The name of the system that logged the message [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Reporter	The name of the user or process that logged the message [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system where the log message originated. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

4.8. General Log Statistics

The General Log Agent is used to collect log messages generated by application programs. The General Log Agent stores these messages in the TeamQuest performance database for analysis. Examples include backup, security, database, and Web server applications.

Table Field Hierarchy

Class:	System
Subclass:	General Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /General Log
TeamQuest Table Name:	System.General Log
Open Table Name:	SYSGENERALLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The message type [Non-Sequential = ID]

4.9. TeamQuest Log Statistics

The following statistics are stored in the performance database tables by the TeamQuest Log Agent. The collection interval and retention periods can be modified. For more information on modifying the collection interval and retention periods, see the *TeamQuest Performance Software Administration Guide*.

Table Field Hierarchy

Class:	Service
Subclass:	TeamQuest Log
IT Resource Name:	/TeamQuest/System/systemname/TeamQuest Log
TeamQuest Table Name:	Service.TeamQuestLog
Open Table Name:	SVCTQLOG
Collection interval:	N/A
Default retention:	1 day
Table type:	Event

Statistic Name	Description
Filename	The name of the TeamQuest log file that was the source of the message text [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The log message type. This is always set to tqlog . [Non-Sequential = ID]

4.10. Derived Statistics

Some products within TeamQuest Performance Software use derived statistics to display common statistics across different platforms. The derived statistics are inserted into the performance database when databases are created. In this section, a derived statistic is marked with an asterisk (*).

You can find information on the following derived statistics:

- Workload Performance Derived Statistics
- TeamQuest On the Web Derived Statistics
- TeamQuest Alert Derived Statistics

4.10.1. Workload Performance Derived Statistics

TeamQuest Manager maintains the following derived statistics that use data from the System Activity Agent and Process-Workload Agent. A derived statistic is marked with an asterisk (*). The workload performance reports reference these statistics. For information on workload performance reports, see the *TeamQuest View Reports Reference Manual*.

Parameter Hierarchy

Class:	Derived
Subclass:	Workload Performance.by Workload
Workload Set:	WLS1, WLS2, ...
Workload:	ALL
Statistic Name:	
%cpu*	The total percentage of CPU utilization. Collected by the Process-Workload Agent.
Kbytes resident memory/process*	The average amount of resident memory used per process. Collected by the Process-Workload Agent.
Kbytes virtual memory/process*	The average amount of virtual memory used per process. Collected by the Process-Workload Agent.
PIOs/sec*	The number of physical I/Os per second. Collected by the Process-Workload Agent.
Population (etime/interval)*	The average number of concurrent processes. Collected by the Process-Workload Agent.
Response (etime/process)*	The elapsed time per process. Collected by the Process-Workload Agent.
Throughput (processes/sec)*	The processes completed per second. Collected by the Process-Workload Agent.
Total Kbytes resident memory*	The average amount of resident memory used. Collected by the Process-Workload Agent.
Total Kbytes virtual memory*	The average amount of virtual memory used. Collected by the Process-Workload Agent.

Class:	Derived
Subclass:	Workload Performance.Summary
Workload Set:	WLS1, WLS2, ...
Statistic Name:	
%cpu*	<p>The total percentage of CPU utilization. Collected by the Process-Workload Agent.</p> <p>View Reports:</p> <p>/report/hp-ux/wkldperf/overall.rpt</p> <p>/report/hp-ux/wkldperf/workload.rpt</p>
block IO avresp*	<p>The average disk I/O response time in milliseconds. Collected by the System Activity Agent.</p> <p>View Report:</p> <p>/report/hp-ux/wkldperf/overall.rpt</p>
block IO r+w/s*	<p>The number of disk I/O response time in milliseconds. Collected by the System Activity Agent.</p> <p>View Report:</p> <p>/report/hp-ux/wkldperf/overall.rpt</p>
Kbytes resident memory/process*	<p>The average amount of resident memory used per process. Collected by the Process-Workload Agent.</p> <p>View Reports:</p> <p>/report/hp-ux/wkldperf/overall.rpt</p> <p>/report/hp-ux/wkldperf/workload.rpt</p>
Kbytes virtual memory/process*	<p>The average amount of virtual memory used per process. Collected by the Process-Workload Agent.</p> <p>View Reports:</p> <p>/report/hp-ux/wkldperf/overall.rpt</p> <p>/report/hp-ux/wkldperf/workload.rpt</p>
PIOs/sec*	<p>The number of physical I/Os per second. Collected by the Process-Workload Agent.</p> <p>View Reports:</p> <p>/report/hp-ux/wkldperf/overall.rpt</p> <p>/report/hp-ux/wkldperf/workload.rpt</p>
Population (etime/interval)*	<p>The average number of concurrent processes. Collected by the Process-Workload Agent.</p> <p>View Reports:</p> <p>/report/hp-ux/wkldperf/overall.rpt</p> <p>/report/hp-ux/wkldperf/workload.rpt</p>
Response (etime/process)*	<p>The elapsed time per process. Collected by the Process-Workload Agent.</p> <p>View Reports:</p> <p>/report/hp-ux/wkldperf/overall.rpt</p> <p>/report/hp-ux/wkldperf/workload.rpt</p>

Throughput (processes/sec)*	The number of processes completed per second. Collected by the Process-Workload Agent. View Reports: /report/hp-ux/wkldperf/overall.rpt /report/hp-ux/wkldperf/workload.rpt
Total Kbytes resident memory*	The average amount of resident memory used for the workload. Collected by the Process-Workload Agent. View Reports: /report/hp-ux/wkldperf/overall.rpt /report/hp-ux/wkldperf/workload.rpt
Total Kbytes virtual memory*	The average amount of virtual memory used for the workload. Collected by the Process-Workload Agent. View Reports: /report/hp-ux/wkldperf/overall.rpt /report/hp-ux/wkldperf/workload.rpt

4.10.2. TeamQuest On the Web Derived Statistics

The derived statistics used by TeamQuest On the Web® include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQWeb.Summary
Statistic Name:	
avg_disk_queue_length*	The average number of requests outstanding for all devices
avg_service_time*	The average amount of time to service each transfer request to all devices in milliseconds
buffer_pct_read_cache*	The percentage of logical reads satisfied from the buffer cache
buffer_pct_write_cache*	The percentage of logical writes satisfied from the buffer cache
disk_xfers_per_sec*	The total number of read and write transfers per second for all devices per second
free_disk_space*	The amount of space available (not in use) on all file systems in megabytes. This measurement is taken at the end of the sampling interval and includes the space held back from normal users.
free_real_mem*	The amount of free memory available in megabytes. This measurement is taken at the end of the sampling interval.
free_swap_space*	The amount of free swap space in megabytes
nfs_calls_per_sec*	The total number of NFS calls sent by the client per second
page_ins_per_sec*	The number of page-in requests per second
page_outs_per_sec*	The number of page-out requests per second

<code>page_scans_per_sec*</code>	The rate per second at which the page daemon scans pages to see if they can be freed
<code>pct_cpu_busy*</code>	The percentage of total CPU time the CPU was busy (not idle). This value includes the time running in system mode and user mode.
<code>pct_disk_busy*</code>	The percentage of time a disk was busy servicing a transfer request
<code>pct_sys_cpu*</code>	The percentage of total CPU time spent running in system mode
<code>pct_usr_cpu*</code>	The percentage of total CPU time spent running in user mode
<code>pkt_errors_per_sec*</code>	The total number (in + out) of network errors per second for all network interfaces
<code>pkts_in_per_sec*</code>	The total number of network input packets per second for all network interfaces
<code>pkts_out_per_sec*</code>	The total number of network output packets per second for all network interfaces
<code>pkts_per_sec*</code>	The total number (in + out) of network packets per second for all network interfaces
<code>total_disk_space*</code>	The total (used + available) amount of space on all file systems in megabytes. This measurement is taken at the end of the sampling interval and includes the space held back from normal users.
<code>total_processes*</code>	The number of entries currently being used in the process table. This measurement is taken at the end of the sampling interval.
<code>total_real_mem*</code>	The total amount of real (physical) memory in megabytes. This measurement is taken at the end of the sampling interval.
<code>total_swap_space*</code>	The total number of megabytes available for swapping

4.10.3. TeamQuest Alert Derived Statistics

The derived statistics used by TeamQuest Alert include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQAlert.Summary
Statistic Name:	
free_real_mem*	The average amount of memory available to user processes in megabytes
net_errors*	The number of network errors for all network interfaces
page_scans*	The rate per second at which the page daemon scans pages to see if they can be freed
pct_swap_free*	The percentage of unused swap space in megabytes at the end of the interval
pct_wio*	The percentage of time spent idle while processes are waiting for I/O completion
run_queue*	The average length of the run queue (a queue of processes in memory and runnable) while the run queue is occupied
total_processes*	The total number of processes active on the system

Section 5

Hyper-V Statistics

The Hyper-V Agent collects performance data on Microsoft Windows Server systems configured in the Hyper-V role. The Hyper-V Agent can collect and store performance information on all Microsoft provided Hyper-V statistics. These statistics include hypervisor, host, and virtual machine related statistics.

This section contains a listing of the statistics collected by the agent:

- Hypervisor Statistics (see 5.1)
- Virtual Machine Statistics (see 5.2)
- Host Statistics (see 5.3)
- I/O Statistics (see 5.4)
- Network Statistics (see 5.5)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

5.1. Hypervisor Statistics

The Hyper-V Agent stores hypervisor statistics in the performance database tables.

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Hypervisor
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Hypervisor
Open Table Name:	HVHYP
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name

Description

Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Logical_Processors	The number of logical processors that are present in the system [Sequential = LST Non-Sequential = SUM]
Monitored_Notifications	The number of monitored notifications that are registered with the hypervisor [Sequential = LST Non-Sequential = SUM]
Partitions	The number of partitions (virtual machines) that are present in the system [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Total_Pages	The number of bootstrap and deposited pages in the hypervisor [Sequential = LST Non-Sequential = SUM]
Virtual_Processors	The number of virtual processors that are present in the system [Sequential = LST Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Hypervisor Logical Processor
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Hypervisor Logical Processor
Open Table Name:	HVHYLOGPROC
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
%_C1_Time	The percentage of time the logical processor spent in the C1 low-power idle state. This statistic is a subset of the total processor idle time. [Sequential = DIV Non-Sequential = AVG]
%_C2_Time	The percentage of time the logical processor spent in the C2 low-power idle state. This statistic is a subset of the total processor idle time. [Sequential = DIV Non-Sequential = AVG]
%_C3_Time	The percentage of time the logical processor spent in the C3 low-power idle state. This statistic is a subset of the total processor idle time. [Sequential = DIV Non-Sequential = AVG]
%_Guest_Run_Time	The percentage of time guest code is running on a logical processor (LP) [Sequential = DIV Non-Sequential = AVG]
%_Hypervisor_Run_Time	The percentage of time the hypervisor is running on a logical processor (LP) [Sequential = DIV Non-Sequential = AVG]
%_Idle_Time	The percentage of time the logical processor (LP) is waiting for work [Sequential = DIV Non-Sequential = AVG]
%_of_Max_Frequency	The percentage of the current processor's maximum frequency. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]

Hyper-V Statistics

%_Total_Run_Time	The percentage of time spent by the logical processor (LP) in guest and hypervisor code. Calculated as $\% \text{ Guest Run Time} + \% \text{ Hypervisor Run Time}$ [Sequential = DIV Non-Sequential = AVG]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
C1_Transitions/sec	The number of times per second the CPU enters the C1 low-power idle state [Sequential = DIV Non-Sequential = SUM]
C2_Transitions/sec	The number of times per second the CPU enters the C2 low-power idle state [Sequential = DIV Non-Sequential = SUM]
C3_Transitions/sec	The number of times per second the CPU enters the C3 low-power idle state [Sequential = DIV Non-Sequential = SUM]
Context_Switches/sec	The number of times per second a new virtual processor (VP) has been scheduled to a logical processor (LP) [Sequential = DIV Non-Sequential = SUM]
Frequency	The frequency of the current processor in megahertz [Sequential = DIV Non-Sequential = AVG]
Hardware_Interrupts/sec	The number of hardware interrupts on the processor (excluding hypervisor interrupts) per second [Sequential = DIV Non-Sequential = SUM]
Inter-Processor_Interrupts_Sent/sec	The number of interrupts sent per second from one processor to another in order to get the processor to do memory coherency [Sequential = DIV Non-Sequential = SUM]
Inter-Processor_Interrupts/sec	The number of inter-processor interrupts (IPI) received per second of a given logical processor (LP) [Sequential = DIV Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Logical_Processor	The logical processor index number. For example, LP 0. [Sequential = ID Non-Sequential = ID]
Monitor_Transition_Cost	The cost to enter the hypervisor through an intercept on a logical processor (LP) [Sequential = DIV Non-Sequential = AVG]
Parking_Status	The parking status of the current processor. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]

Processor_State_Flags	The state flags of the current processor. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
Root_VP_Index	The index of the root virtual processor that is affinity bound to the logical processor. A value greater than the maximum possible root virtual processor (VP) index indicates no binding. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Scheduler_Interrupts/sec	The number of interrupts sent per second by the hypervisor scheduler from one logical processor (LP) to another [Sequential = DIV Non-Sequential = SUM]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Timer_Interrupts/sec	The number of hypervisor timer interrupts per second on the logical processor (LP) [Sequential = DIV Non-Sequential = SUM]
Total_Interrupts/sec	The number of hardware and hypervisor interrupts per second [Sequential = DIV Non-Sequential = SUM]

Hyper-V Statistics

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Hypervisor Partition
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.Hypervisor Partition
Open Table Name:	HVHYPPART
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
1_G_device_pages	The number of 1 GB pages present in the device space of the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
1G_GPA_pages	The number of 1 GB pages present in the guest physical address (GPA) space of the partition [Sequential = LST Non-Sequential = SUM]
2_M_device_pages	The number of 2 MB pages present in the device space of the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
2M_GPA_pages	The number of 2 MB pages present in the guest physical address (GPA) space of the partition [Sequential = LST Non-Sequential = SUM]
4_K_device_pages	The number of 4 KB pages present in the device space of the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
4K_GPA_pages	The number of 4 KB pages present in the guest physical address (GPA) space of the partition [Sequential = LST Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Address_Spaces	The number of address spaces in the virtual translation lookaside buffer (TLB) of the partition [Sequential = LST Non-Sequential = SUM]

Attached_Devices	The number of devices attached to the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Deposited_Pages	The number of pages that are deposited into the partition. This value shows how much memory the hypervisor is using for managing the partition. [Sequential = LST Non-Sequential = SUM]
Device_DMA_Errors	The number of illegal direct memory access requests generated by all devices assigned to the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Device_Interrupt_Errors	The number of illegal interrupt requests generated by all devices assigned to the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Device_Interrupt_Mappings	The number of device interrupt mappings used by the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Device_Interrupt_Throttle_Events	The number of times an interrupt from a device assigned to the partition was temporarily throttled because the device was generating too many interrupts. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
GPA_Pages	The number of pages that are present in the guest physical address (GPA) space of the partition. This value is 0 for the host. [Sequential = LST Non-Sequential = SUM]
GPA_Space_Modifications/sec	The number of modifications to the guest physical address (GPA) space of the partition per second [Sequential = DIV Non-Sequential = SUM]
I/O_TLB_Flush_Cost	The average amount of time in nanoseconds spent processing an I/O translation lookaside buffer (TLB) flush. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
I/O_TLB_Flushes/sec	The number of flushes of I/O translation lookaside buffers (TLB) of the partition per second. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Recommended_Virtual_TLB_Size	The recommended number of pages to be deposited for the virtual translation lookaside buffer (TLB) [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

Hyper-V Statistics

Skipped_Timer_Ticks	The number of timer interrupts skipped for the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]
Virtual_Processors	The number of virtual processors that are present in the partition [Sequential = LST Non-Sequential = SUM]
Virtual_TLB_Flush_Entries/sec	The number of flushes per second on the virtual translation lookaside buffer (TLB) [Sequential = AVG Non-Sequential = SUM]
Virtual_TLB_Pages	The number of pages used by the virtual translation lookaside buffer (TLB) [Sequential = LST Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Hypervisor Root Partition
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Hypervisor Root Partition
Open Table Name:	HVHYPROOTPART
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Type	Description
1_G_device_pages	Double	The number of 1 GB pages present in the device space of the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
1G_GPA_pages	Double	The number of 1 GB pages present in the guest physical address (GPA) space of the partition [Sequential = LST Non-Sequential = SUM]

2_M_device_pages	Double	The number of 2 MG pages present in the device space of the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
2M_GPA_pages	Double	The number of 2 MB pages present in the guest physical address (GPA) space of the partition [Sequential = LST Non-Sequential = SUM]
4_K_device_pages	Double	The number of 4 KB pages present in the device space of the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
4K_GPA_pages	Double	The number of 4 KB pages present in the guest physical address (GPA) space of the partition [Sequential = LST Non-Sequential = SUM]
Actual_Interval	Long	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Address_Spaces	Double	The number of address spaces in the virtual translation lookaside buffer (TLB) of the partition [Sequential = LST Non-Sequential = SUM]
Attached_Devices	Double	The number of devices attached to the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Deposited_Pages	Double	The number of pages that are deposited into the partition. This value shows how much memory the hypervisor is using for managing the partition. [Sequential = LST Non-Sequential = SUM]
Device_DMA_Errors	Double	The number of illegal direct memory access requests generated by all devices assigned to the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Device_Interrupt_Errors	Double	The number of illegal interrupt requests generated by all devices assigned to the partition [Sequential = LST Non-Sequential = SUM]

Hyper-V Statistics

Device_Interrupt_Mappings	Double	The number of device interrupt mappings used by the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Device_Interrupt_Throttle_Events	Double	The number of times an interrupt from a device assigned to the partition was temporarily throttled because the device was generating too many interrupts. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
GPA_Pages	Double	The number of pages that are present in the guest physical address (GPA) space of the partition. This value is 0 for the host. [Sequential = LST Non-Sequential = SUM]
GPA_Space_Modifications/sec	Double	The number of modifications to the guest physical address (GPA) space of the partition per second [Sequential = DIV Non-Sequential = SUM]
I/O_TLB_Flush_Cost	Double	The average amount of time in nanoseconds spent processing an I/O translation lookaside buffer (TLB) flush [Sequential = AVG Non-Sequential = SUM]
I/O_TLB_Flushes/sec	Double	The number of flushes of I/O translation lookaside buffers (TLB) of the partition per second [Sequential = AVG Non-Sequential = SUM]
Interval	Long	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Partition	String	The name of the Hyper-V host [Sequential = ID Non-Sequential = ID]
Recommended_Virtual_TLB_Size	Double	The recommended number of pages to be deposited for the virtual translation lookaside buffer (TLB) [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	Double	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Skipped_Timer_Ticks	Double	The number of timer interrupts skipped for the partition. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
System	String	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]

Time	Long	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Processors	Double	The number of virtual processors that are present in the partition [Sequential = LST Non-Sequential = SUM]
Virtual_TLB_Flush_Entires/sec	Double	The number of flushes on the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Virtual_TLB_Pages	Double	The number of pages used by the virtual translation lookaside buffer (TLB) [Sequential = LST Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Hypervisor Root Virtual Processor
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Hypervisor Root Virtual Processor
Open Table Name:	HVHYPROOTVPROC
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
%_Guest_Run_Time	The percentage of time the root virtual processor (VP) is running in non-hypervisor code on a logical processor (LP) [Sequential = DIV Non-Sequential = AVG]
%_Hypervisor_Run_Time	The percentage of time the root virtual processor (VP) is running in hypervisor code on a logical processor (LP) [Sequential = DIV Non-Sequential = AVG]
%_Remote_Run_Time	The percentage of time spent by the virtual processor in hypervisor code. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
%_Total_Run_Time	The percentage of time spent by the virtual processor in guest and hypervisor code per virtual processor. Calculated as % Guest Run Time + % Hypervisor Run Time [Sequential = DIV Non-Sequential = AVG]

Hyper-V Statistics

Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Address_Domain_Flushes/sec	The number of explicit flushes of the virtual translation lookaside buffer (TLB) by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Address_Space_Evictions/sec	The number of address space evictions in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Address_Space_Flushes/sec	The number of explicit flushes of one address space by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Address_Space_Switches/sec	The number of address space switches by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
APIC_EOI_Accesses/sec	The number of Advanced Programmable Interrupt Controller (APIC) end of interrupt (EOI) register writes by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
APIC_IPIs_Sent/sec	The number of Advanced Programmable Interrupt Controller (APIC) inter-processor interrupts (IPI) (including to self) sent to the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
APIC_MMIO_Accesses/sec	The number of Advanced Programmable Interrupt Controller (APIC) Memory-mapped I/O (MMIO) register accesses by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
APIC_Self_IPIs_Sent/sec	The number of Advanced Programmable Interrupt Controller (APIC) inter-processor interrupts (IPI) sent by the virtual processor to itself per second [Sequential = DIV Non-Sequential = SUM]
APIC_TPR_Accesses/sec	The number of Advanced Programmable Interrupt Controller (APIC) task priority (TPR) accesses by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Control_Register_Accesses/sec	The number of control register accesses by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Control_Register_Accesses_Cost	The average amount of time in nanoseconds spent processing a control register access per second [Sequential = DIV Non-Sequential = AVG]

CPU_Wait_Time_Per_Dispatch	The average amount of time in nanoseconds spent waiting for a virtual processor to be dispatched onto a logical processor. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
CPUID_Instructions_Cost	The relative measure of the cost of the CPUID instruction. The CPUID instruction is used to retrieve information on the local CPU capabilities. [Sequential = DIV Non-Sequential = AVG]
CPUID_Instructions/sec	The number of CPUID instruction calls per second [Sequential = DIV Non-Sequential = SUM]
Debug_Register_Accesses_Cost	The average amount of time in nanoseconds spent handling a debug register access [Sequential = DIV Non-Sequential = AVG]
Debug_Register_Accesses/sec	The number of debug register accesses by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Emulated_Instructions_Cost	The relative measure of the cost of emulation [Sequential = DIV Non-Sequential = AVG]
Emulated_Instructions/sec	The number of emulated instructions completed per second [Sequential = DIV Non-Sequential = SUM]
External_Interrupts_Cost	The average amount of time in nanoseconds spent processing an external interrupt [Sequential = DIV Non-Sequential = AVG]
External_Interrupts/sec	The number of external interrupts received by the hypervisor while executing guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Global_GVA_Range_Flushes/sec	The number of explicit flushes of a guest virtual address (GVA) range in all address spaces by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
GPA_Space_Hypercalls/sec	The number of guest physical address (GPA) space hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Guest_Page_Table_Maps/sec	The number of map operations for guest page table pages per second [Sequential = DIV Non-Sequential = SUM]
Hardware_Interrupts/sec	The number of hardware interrupts from attached devices on the virtual processor. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
HLT_Instructions_Cost	The average amount of time in nanoseconds spent processing a halt (HLT) instruction [Sequential = DIV Non-Sequential = AVG]

Hyper-V Statistics

HLT_Instructions/sec	The number of CPU halts per second on the virtual processor (VP) [Sequential = DIV Non-Sequential = SUM]
Hypercalls_Cost	The average amount of time in nanoseconds spent processing a hypercall [Sequential = DIV Non-Sequential = AVG]
Hypercalls/sec	The number of hypercalls made by the guest code on the virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = DIV Non-Sequential = AVG]
IO_Instructions_Cost	The average amount of time in nanoseconds spent processing an I/O instruction [Sequential = DIV Non-Sequential = AVG]
IO_Instructions/sec	The number of I/O instructions executed by guest code on a virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
IO_Intercept_Messages/sec	The number of I/O intercept messages sent to the parent partition per second [Sequential = DIV Non-Sequential = SUM]
Large_Page_TLB_Fills/sec	The number of large page translation lookaside buffer (TLB) fills per second [Sequential = DIV Non-Sequential = SUM]
Local_Flushed_GVA_Ranges/sec	The number of explicit flushes of a guest virtual address (GVA) range in one address space by guest code on the virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
Logical_Processor_Dispatches/sec	The number of dispatches of this virtual processor onto logical processors. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Logical_Processor_Hypercalls/sec	The number of logical processor hypercalls made by guest code on the virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
Logical_Processor_Migrations/sec	The number of migrations by the virtual processor to a logical processor per second [Sequential = DIV Non-Sequential = SUM]
Long_Spin_Wait_Hypercalls/sec	The number of long spin wait hypercalls made by guest code on the virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
Memory_Intercept_Messages/sec	The number of memory intercept messages sent to the parent partition per second [Sequential = DIV Non-Sequential = SUM]
MSR_Accesses_Cost	The relative measure of the cost of the machine-specific register (MSR) instruction [Sequential = DIV Non-Sequential = AVG]

MSR_Accesses/sec	The number of machine-specific register (MSR) instruction calls per second [Sequential = DIV Non-Sequential = SUM]
MWAIT_Instructions_Cost	The relative measure of the cost of the MONITOR WAIT (MWAIT) instructions [Sequential = DIV Non-Sequential = AVG]
MWAIT_Instructions/sec	The number of MONITOR WAIT (MWAIT) instructions per second [Sequential = DIV Non-Sequential = SUM]
Nested_Page_Fault_Intercepts_Cost	The average amount of time in nanoseconds spent processing a nested page fault intercept. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
Nested_Page_Fault_Intercepts/sec	The number of nested page fault exceptions intercepted by the hypervisor while executing the guest virtual processor. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
Other_Hypercalls/sec	The number of other hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Other_Intercepts_Cost	The average amount of time in nanoseconds spent processing other intercepts [Sequential = DIV Non-Sequential = AVG]
Other_Intercepts/sec	The number of other intercepts triggered by guest code on a virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Other_Messages/sec	The number of other intercept messages sent to the parent partition per second [Sequential = DIV Non-Sequential = SUM]
Page_Fault_Intercepts_Cost	The relative measure of the cost of a page fault intercept [Sequential = DIV Non-Sequential = AVG]
Page_Fault_Intercepts/sec	The number of page fault intercepts per second. Whenever guest code accesses a page not in the CPU translation lookaside buffer (TLB), a page fault occurs. [Sequential = DIV Non-Sequential = SUM]
Page_Invalidations_Cost	The average amount of time in nanoseconds spent processing an invalid page (INVLPG) instruction [Sequential = DIV Non-Sequential = AVG]
Page_Invalidations/sec	The number of invalid page (INVLPG) instructions executed by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Allocations/sec	The number of page table allocations in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]

Hyper-V Statistics

Page_Table_Evictions/sec	The number of page table evictions in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Reclamations/sec	The number of reclamations of unreferenced page tables in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Resets/sec	The number of page table resets in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Validations/sec	The number of page table validations to remove state entries in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Write_Intercepts/sec	The number of write intercepts on guest page tables by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Pending_Interrupts_Cost	The average amount of time in nanoseconds spent processing a pending interrupt intercept [Sequential = DIV Non-Sequential = AVG]
Pending_Interrupts/sec	The number of interrupts per second due to a task priority (TPR) reduction by guest code on the virtual processor [Sequential = DIV Non-Sequential = SUM]
Reflected_Guest_Page_Faults/sec	The number of page fault exceptions delivered to the guest per second [Sequential = DIV Non-Sequential = SUM]
Root_Virtual_Processor	The root virtual processor index number. For example, VP 0. [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Small_Page_TLB_Fills/sec	The number of virtual translation lookaside buffer (TLB) misses on 4K pages per second [Sequential = DIV Non-Sequential = SUM]
Synthetic_Interrupt_Hypercalls/sec	The number of synthetic interrupt hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Synthetic_Interrupts/sec	The number of synthetic interrupts delivered to the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Total_Intercepts_Cost	The relative measure of the cost of intercepts [Sequential = DIV Non-Sequential = AVG]
Total_Intercepts/sec	The number of intercepts per second. An intercept occurs whenever a guest virtual processor (VP) needs to exit its current mode of running for servicing in the hypervisor. [Sequential = DIV Non-Sequential = SUM]
Total_Messages/sec	The number of messages sent to the parent partition per second [Sequential = DIV Non-Sequential = SUM]
Virtual_Interrupt_Hypercalls/sec	The number of virtual interrupt hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Virtual_Interrupts/sec	The number of interrupts (including synthetic interrupts) delivered to the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Virtual_MMU_Hypercalls/sec	The number of virtual memory management unit (MMU) hypercalls made by the guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Virtual_Processor_Hypercalls/sec	The number of virtual processor hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Hypervisor Virtual Processor
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.Hypervisor Virtual Processor
Open Table Name:	HVHYPVPROC
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
%_Guest_Run_Time	The percentage of time the guest virtual processor (VP) is running in non-hypervisor code on a logical processor (LP) [Sequential = DIV Non-Sequential = AVG]
%_Hypervisor_Run_Time	The percentage of time the guest virtual processor (VP) is running in hypervisor code on a logical processor (LP) [Sequential = DIV Non-Sequential = AVG]

Hyper-V Statistics

%_Remote_Run_Time	The percentage of time spent by the virtual processor in hypervisor code. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
%_Total_Run_Time	The percentage of time spent by the virtual processor in guest and hypervisor code per virtual processor. Calculated as % Guest Run Time + % Hypervisor Run Time [Sequential = DIV Non-Sequential = AVG]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Address_Domain_Flushes/sec	The number of explicit flushes of the virtual translation lookaside buffer (TLB) by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Address_Space_Evictions/sec	The number of address space evictions in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Address_Space_Flushes/sec	The number of explicit flushes of one address space by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Address_Space_Switches/sec	The number of address space switches by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
APIC_EOI_Accesses/sec	The number of Advanced Programmable Interrupt Controller (APIC) end of interrupt (EOI) register writes by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
APIC_IPIs_Sent/sec	The number of Advanced Programmable Interrupt Controller (APIC) inter-processor interrupts (IPI) (including to self) sent to the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
APIC_MMIO_Accesses/sec	The number of Advanced Programmable Interrupt Controller (APIC) Memory-mapped I/O (MMIO) register accesses by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
APIC_Self_IPIs_Sent/sec	The number of Advanced Programmable Interrupt Controller (APIC) inter-processor interrupts (IPI) sent by the virtual processor to itself per second [Sequential = DIV Non-Sequential = SUM]
APIC_TPR_Accesses/sec	The number of Advanced Programmable Interrupt Controller (APIC) task priority (TPR) accesses by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]

Control_Register_Accesses/sec	The number of control register accesses by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Control_Register_Accesses_Cost	The average amount of time in nanoseconds spent processing a control register access per second [Sequential = DIV Non-Sequential = AVG]
CPU Wait Time Per Dispatch	The average amount of time in nanoseconds spent waiting for a virtual processor to be dispatched onto a logical processor. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
CPUID_Instructions_Cost	The relative measure of the cost of the CPUID instruction. CPUID instruction is used to retrieve information on the local CPU capabilities. [Sequential = DIV Non-Sequential = AVG]
CPUID_Instructions/sec	The number of CPUID instruction calls per second [Sequential = DIV Non-Sequential = SUM]
Debug_Register_Accesses_Cost	The average amount of time in nanoseconds spent handling a debug register access [Sequential = DIV Non-Sequential = AVG]
Debug_Register_Accesses/sec	The number of debug register accesses by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Emulated_Instructions_Cost	The relative measure of the cost of instruction emulation [Sequential = DIV Non-Sequential = AVG]
Emulated_Instructions/sec	The number of emulated instructions completed per second [Sequential = DIV Non-Sequential = SUM]
External_Interrupts_Cost	The average amount of time in nanoseconds spent processing an external interrupt [Sequential = DIV Non-Sequential = AVG]
External_Interrupts/sec	The number of external interrupts received by the hypervisor while executing guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Global_GVA_Range_Flushes/sec	The number of explicit flushes of a guest virtual address (GVA) range in all address spaces by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
GPA_Space_Hypercalls/sec	The number of guest physical address (GPA) space hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Guest_Page_Table_Maps/sec	The number of map operations for guest page table pages per second [Sequential = DIV Non-Sequential = SUM]
Hardware_Interrupts/sec	The number of hardware interrupts from attached devices on the virtual processor. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]

Hyper-V Statistics

HLT_Instructions_Cost	The average amount of time in nanoseconds spent processing a halt (HLT) instruction [Sequential = DIV Non-Sequential = AVG]
HLT_Instructions/sec	The number of CPU halts per second on the virtual processor (VP) [Sequential = DIV Non-Sequential = SUM]
Hypercalls_Cost	The average amount of time in nanoseconds spent processing a hypercall [Sequential = DIV Non-Sequential = AVG]
Hypercalls/sec	The number of hypercalls made by the guest code on the virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
IO_Instructions_Cost	The average amount of time in nanoseconds spent processing an I/O instruction [Sequential = DIV Non-Sequential = AVG]
IO_Instructions/sec	The number of I/O instructions executed by guest code on a virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
IO_Intercept_Messages/sec	The number of I/O intercept messages sent to the parent partition per second [Sequential = DIV Non-Sequential = SUM]
Large_Page_TLB_Fills/sec	The number of large page translation lookaside buffer (TLB) fills per second [Sequential = DIV Non-Sequential = SUM]
Local_Flushed_GVA_Ranges/sec	The number of explicit flushes of a guest virtual address (GVA) range in one address space by guest code on the virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
Logical_Processor_Dispatches/sec	The number of dispatches of the virtual processor onto logical processors. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
Logical_Processor_Hypercalls/sec	The number of logical processor hypercalls made by guest code on the virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
Logical_Processor_Migrations/sec	The number of migrations by the virtual processor to a logical processor per second [Sequential = DIV Non-Sequential = SUM]
Long_Spin_Wait_Hypercalls/sec	The number of long spin wait hypercalls made by guest code on the virtual processor (VP) per second [Sequential = DIV Non-Sequential = SUM]
Memory_Intercept_Messages/sec	The number of memory intercept messages sent to the parent partition per second [Sequential = DIV Non-Sequential = SUM]

MSR_Accesses_Cost	The relative measure of the cost of the machine-specific register (MSR) instruction [Sequential = DIV Non-Sequential = AVG]
MSR_Accesses/sec	The number of machine-specific register (MSR) instruction calls per second [Sequential = DIV Non-Sequential = SUM]
MWAIT_Instructions_Cost	The relative measure of the cost of the MONITOR WAIT (MWAIT) instructions [Sequential = DIV Non-Sequential = AVG]
MWAIT_Instructions/sec	The number of MONITOR WAIT (MWAIT) instructions per second [Sequential = DIV Non-Sequential = SUM]
Nested_Page_Fault_Intercepts_Cost	The average amount of time in nanoseconds spent processing a nested page fault intercept. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
Nested_Page_Fault_Intercepts/sec	The number of nested page fault exceptions intercepted by the hypervisor while executing the guest virtual processor. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = AVG]
Other_Hypercalls/sec	The number of other hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Other_Intercepts_Cost	The average amount of time in nanoseconds spent processing other intercepts [Sequential = DIV Non-Sequential = AVG]
Other_Intercepts/sec	The number of other intercepts triggered by guest code on a virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Other_Messages/sec	The number of other intercept messages sent to the parent partition per second [Sequential = DIV Non-Sequential = SUM]
Page_Fault_Intercepts_Cost	The relative measure of the cost of a page fault intercept [Sequential = DIV Non-Sequential = AVG]
Page_Fault_Intercepts/sec	The number of page fault intercepts per second. Whenever guest code accesses a page not in the CPU translation lookaside buffer (TLB), a page fault occurs. [Sequential = DIV Non-Sequential = SUM]
Page_Invalidations_Cost	The average amount of time in nanoseconds spent processing an invalid page (INVLPG) instruction [Sequential = DIV Non-Sequential = AVG]
Page_Invalidations/sec	The number of invalid page (INVLPG) instructions executed by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Allocations/sec	The number of page table allocations in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]

Hyper-V Statistics

Page_Table_Evictions/sec	The number of page table evictions in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Reclamations/sec	The number of reclamations of unreferenced page tables in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Resets/sec	The number of page table resets in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Validations/sec	The number of page table validations to remove state entries in the virtual translation lookaside buffer (TLB) per second [Sequential = DIV Non-Sequential = SUM]
Page_Table_Write_Intercepts/sec	The number of write intercepts on guest page tables by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Pending_Interrupts_Cost	The average amount of time in nanoseconds spent processing a pending interrupt intercept [Sequential = DIV Non-Sequential = AVG]
Pending_Interrupts/sec	The number of interrupts per second due to a task priority (TPR) reduction by guest code on the virtual processor [Sequential = DIV Non-Sequential = SUM]
Reflected_Guest_Page_Faults/sec	The number of page fault exceptions delivered to the guest per second [Sequential = DIV Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Small_Page_TLB_Fills/sec	The number of virtual translation lookaside buffer (TLB) misses on 4K pages per second [Sequential = DIV Non-Sequential = SUM]
Synthetic_Interrupt_Hypercalls/sec	The number of synthetic interrupt hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Synthetic_Interrupts/sec	The number of synthetic interrupts delivered to the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Total_Intercepts_Cost	The relative measure of the cost of intercepts [Sequential = DIV Non-Sequential = AVG]
Total_Intercepts/sec	The number of intercepts per second. An intercept occurs whenever a guest virtual processor (VP) needs to exit its current mode of running for servicing in the hypervisor. [Sequential = DIV Non-Sequential = SUM]

Total_Messages/sec	The number of messages sent to the parent partition per second [Sequential = DIV Non-Sequential = SUM]
Virtual_Interrupt_ Hypercalls/sec	The number of virtual interrupt hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Virtual_Interrupts/sec	The number of interrupts (including synthetic interrupts) delivered to the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]
Virtual_MMU_ Hypercalls/sec	The number of virtual memory management unit (MMU) hypercalls made by the guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Virtual_Processor_ Hypercalls/sec	The number of virtual processor hypercalls made by guest code on the virtual processor per second [Sequential = DIV Non-Sequential = SUM]
Virtual_Processor	The virtual processor index number. For example, VP 0 [Sequential = ID Non-Sequential = ID]

5.2. Virtual Machine Statistics

The Hyper-V Agent stores virtual machine statistics in the performance database tables.

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Dynamic Memory VM
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V. Dynamic Memory VM
Open Table Name:	HVDYNMEMVM
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Added_Memory	The amount of memory added to the virtual machine (VM). This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Average_Pressure	The average pressure in the virtual machine (VM) [Sequential = LST Non-Sequential = AVG]
Current_Pressure	The current pressure in the virtual machine (VM) [Sequential = LST Non-Sequential = SUM]
Guest_Visible_Physical_Memory	The amount of memory visible in the virtual machine (VM) [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Maximum_Pressure	The maximum pressure band in the virtual machine (VM) [Sequential = LST Non-Sequential = MAX]
Memory_Add_Operations	The number of add operations for the virtual machine (VM). This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Memory_Remove_Operations	The number of remove operations for the virtual machine (VM). This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]

Minimum_Pressure	The minimum pressure band in the virtual machine (VM) [Sequential = LST Non-Sequential = MIN]
Physical_Memory	The current amount of memory in the virtual machine (VM) [Sequential = LST Non-Sequential = SUM]
Removed_Memory	The amount of memory removed from the virtual machine (VM). This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Smart_Paging_Working_Set_Size	The size of the memory in megabytes backed by smart paging. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]

Hyper-V Statistics

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Replica VM
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.Replica VM
Open Table Name:	HVREPLICAVM
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Average Replication Latency	The average replication latency in seconds [Sequential = AVG Non-Sequential = AVG]
Average Replication Size	The average replication size in bytes [Sequential = AVG Non-Sequential = AVG]
Compression Efficiency	The compression efficiency for the files that have been transferred over the network [Sequential = AVG Non-Sequential = AVG]
Last Replication Size	The last replication size in bytes [Sequential = LST Non-Sequential = SUM]
Network Bytes Recv	The total bytes received over the network for the virtual machines since the virtual machine management service was started [Sequential = AVG Non-Sequential = AVG]
Network Bytes Sent	The total bytes sent over the network for the virtual machines since the virtual machine management service was started [Sequential = AVG Non-Sequential = AVG]
Replication Count	The replication count since the virtual machine management service was started [Sequential = SUM Non-Sequential = SUM]
Replication Latency	The last replication latency in seconds. This value represents the time taken for the delta to be applied on the recovery since it was snapped. [Sequential = LST Non-Sequential = SUM]
Resynchronized Bytes	The total bytes sent and received over the network for the virtual machine during the resynchronize operation since the virtual machine management service was started [Sequential = SUM Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual Machines
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Virtual Machines
Open Table Name:	HVVMS
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
EnabledState	<p>The enabled status of the virtual machine. The enabled state can be any of the following:</p> <ul style="list-style-type: none">Disabled = the virtual machine is turned offEnabled = the virtual machine is runningPaused = the virtual machine is not runningPausing = the virtual machine is transitioning from Enabled to PausedResuming = the virtual machine is transitioning from Paused to Enabled and is resuming from a paused stateSaving = the virtual machine is transitioning from Enabled to Suspended and is saving its stateSnapshotting = the virtual machine is performing a snapshot operationStarting = the virtual machine is transitioning from Disabled or Suspended to EnabledStopping = the virtual machine is transitioning from Enabled to Disabled and is turning offSuspended = the virtual machine is in a saved stateUnknown = the virtual machine state cannot be determined <p>[Non-Sequential = ID]</p>
GuestOperatingSystem	<p>The name of the guest operating system or an empty string if none exists</p> <p>[Non-Sequential = ID]</p>

HealthState	<p>The health status of the virtual machine. The health status can be any of the following:</p> <p>Critical Failure = the process for the virtual machine (vmwp.exe) is not responding</p> <p>Degraded/Warning = the virtual machine is not operating at optimal performance or is reporting recoverable errors</p> <p>Major Failure = the virtual hard disks are low on disk space and the virtual machine has been paused</p> <p>Minor Failure = all functionality of the virtual machine is available but some functionality might be degraded</p> <p>Non-Recoverable Error = the virtual machine completely failed and recovery is not possible</p> <p>OK = the virtual machine is functional and operating normally</p> <p>Unknown = the virtual machine health state cannot be determined</p> <p>[Non-Sequential = ID]</p>
Heartbeat	<p>The heartbeat status of the virtual machine. The heartbeat status can be any of the following:</p> <p>Degraded = the virtual machine service negotiated a compatible communications protocol version</p> <p>Lost Communication = the virtual machine service is not responding</p> <p>No Contact = the virtual machine service cannot be contacted</p> <p>Non-Recoverable Error = the service does not support a compatible protocol version</p> <p>Unavailable = the host does not expect a heartbeat from the virtual machine</p> <p>Unknown = the virtual machine heartbeat cannot be determined</p> <p>[Non-Sequential = ID]</p>
Limit	<p>The percentage of the limit of the host processor for the virtual machine. The Limit status controls the maximum amount of processing the virtual machine can receive. This field comes from Hyper-V Manager > Virtual Machines > Settings > Processor > Resource Control > Virtual machine limit (percentage) field.</p> <p>[Non-Sequential = AVG]</p>
MemoryType	<p>The type of memory the virtual machine is using (static or dynamic)</p> <p>[Non-Sequential = ID]</p>
MemoryUsage	<p>The current memory usage of the virtual machine in megabytes (MB)</p> <p>[Non-Sequential = AVG]</p>
Notes	<p>The notes associated with the virtual machine</p> <p>[Non-Sequential = ID]</p>
NumberOfProcessors	<p>The number of processors allocated to the virtual machine</p> <p>[Non-Sequential = AVG]</p>

ProcessID	The unique identifier of the instance of Virtual Machine Work Process (vmwp.exe) instance running on the Hyper-V host that represents the virtual machine (static or dynamic) [Non-Sequential = ID]
Reservation	The percentage of the reservation for the host processor for the virtual machine. The Reservation status controls the amount of processing reserved for the virtual machine. This field comes from Hyper-V Manager > Virtual Machines > Settings... > Processor > Resource Control > Virtual machine reserve (percentage) field. [Non-Sequential = AVG]
Status	<p>The operational status of the virtual machine. On Microsoft Windows systems, this field is OperationalStatus. The status can be any of the following:</p> <p>Applying Snapshot = the virtual machine is applying a snapshot</p> <p>Creating Snapshot = the virtual machine is creating a snapshot</p> <p>Degraded = the virtual machine storage containing the configuration is not accessible</p> <p>Deleting Snapshot = the virtual machine is deleting a snapshot</p> <p>Dormant = the virtual machine is suspended or paused</p> <p>Exporting Virtual Machine = the virtual machine is being exported</p> <p>In Service = the virtual machine is processing a request</p> <p>Merging Disks = the virtual hard disks from deleted snapshots are being merged</p> <p>Migrating Virtual Machine = the virtual machine is migrating from one physical system to another</p> <p>OK = the virtual machine is functional and operating normally</p> <p>Predictive Failure = the virtual hard disks are low on free space</p> <p>Stopped = the virtual machine is no longer running</p> <p>Waiting to Start = the virtual machine will start after the AutomaticStartupActionDelay has elapsed</p> <p>[Non-Sequential = ID]</p>
System	The name of the Hyper-V system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
TimeOfLastConfigurationChange	The date and time of the last modification of the virtual machine configuration file. On Microsoft Windows systems, this field is TimeOfLastConfigurationChange. [Non-Sequential = LST]
TimeOfLastStateChange	The date and time of the last virtual machine state change. On Microsoft Windows systems, this field is TimeOfLastStateChange. [Non-Sequential = LST]

Hyper-V Statistics

UpTime	The amount of time since the virtual machine was last booted [Non-Sequential = LST]
VirtualMachine	The name of the virtual machine. On Microsoft Windows systems, this field is ElementName. [Non-Sequential = ID]
Weight	The host processor weight for the virtual machine. The Weight status controls the amount of processing the virtual machine receives compared to the other virtual machines. If the value is higher than other virtual machine weight, this virtual machine should receive more processor time. This field comes from the Hyper-V Manager > Virtual Machines > Settings... > Processor > Resource Control > Relative weight field. [Non-Sequential = AVG]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	VM Virtual Network Adapter
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.VM Virtual Network Adapter
Open Table Name:	HVVMVIRTNETADAPTER
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Broadcast_Packets_Received/sec	The number of broadcast packets received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Broadcast_Packets_Sent/sec	The number of broadcast packets sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Bytes_Received/sec	The number of bytes received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Bytes_Sent/sec	The number of bytes sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]

Bytes/sec	The number of bytes sent or received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Directed_Packets_Received/sec	The number of directed packets received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Directed_Packets_Sent/sec	The number of directed packets sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Dropped_Packets_Incoming/sec	The total number of incoming packets dropped per second by the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Dropped_Packets_Outgoing/sec	The total number of outgoing packets dropped per second by the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Extensions_Dropped_Packets_Incoming/sec	The total number of incoming packets dropped per second by switch extensions of the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Extensions_Dropped_Packets_Outgoing/sec	The total number of outgoing packets dropped per second by switch extensions of the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
GUID	The globally unique identifier (GUID) of the virtual machine virtual network adapter [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
IPsec_offload_Bytes_Receive/sec	The total number of Psec offload bytes received per second by the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
IPsec_offload_Bytes_Sent/sec	The total number of Psec offload bytes sent per second by the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Multicast_Packets_Received/sec	The number of multicast packets received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Multicast_Packets_Sent/sec	The number of multicast packets sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Packets_Received/sec	The number of packets received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]

Hyper-V Statistics

Packets_Sent/sec	The number of packets sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Packets/sec	The number of packets sent or received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	VM Save, Snapshot, and Restore
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.VM Save, Snapshot, and Restore
Open Table Name:	HVVMSAVESNAPSHOT
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Buffers_Compressed_on_I/O_Thread	The number of RAM buffers compressed on the disk I/O thread [Sequential = LAST Non-Sequential = SUM]
Buffers_Saved	The total number of RAM buffers processed during a save or snapshot operation [Sequential = LAST Non-Sequential = SUM]
Operation_Time	The amount of time taken to complete a Global Memory Optimizer (GMO) operation in milliseconds. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Requests_Active	The total number of requests currently being processed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]

Requests_Dispatched	The total number of requests that have been dispatched. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Requests_High_Priority	The total number of high priority requests. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Requests_Processed	The total number of requests that have been processed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Save_Time	The amount of time it takes to complete a save operation in milliseconds [Sequential = LAST Non-Sequential = SUM]
Snapshot_Buffer_Copy_on_Intercepts	The number of RAM buffers copied as a result of guest access during an ongoing snapshot [Sequential = LAST Non-Sequential = SUM]
Snapshot_Time	The amount of time it takes to complete a snapshot operation in milliseconds [Sequential = LAST Non-Sequential = SUM]
Threads_Spawned	The total number of threads currently spawned. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]

Hyper-V Statistics

Table Field Hierarchy

Class:	Hyper-V
Subclass:	VM Remoting
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.VM Remoting
Open Table Name:	HVVMREMOTING
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Connected_Clients	The number of connected clients to the virtual machine (VM) [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Updated_Pixels/sec	The number of pixels that are updated per second for the virtual machine (VM) [Sequential = AVG Non-Sequential = AVG]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	VM Vid Partition
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.VM Vid Partition
Open Table Name:	HVVMVIDPART
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Physical_Pages_ Allocated	The number of physical pages allocated [Sequential = LST Non-Sequential = SUM]
Preferred_NUMA_Node	The preferred Non-Uniform Memory Access (NUMA) node associated with the partition [Sequential = ID Non-Sequential = ID]
Remote_Physical_Pages	The number of physical pages not allocated from the preferred Non-Uniform Memory Access (NUMA) node [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]

5.3. Host Statistics

The Hyper-V Agent stores host statistics in the performance database tables.

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Dynamic Memory Integration Service
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Dynamic Memory Integration Service
Open Table Name:	HVDYNMEMINT
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Maximum Memory, Mbytes	The maximum amount of memory in megabytes available to the Dynamic Memory Balancer [Sequential = DIV Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Dynamic Memory Balancer
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Dynamic Memory Balancer
Open Table Name:	HVDYNMEMBALANCER
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Added_Memory	The amount of memory added to virtual machines (VM). This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Available_Memory	The amount of memory left on the node [Sequential = LST Non-Sequential = SUM]
Average_Pressure	The average pressure on the balancer node [Sequential = LST Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Memory_Add_Operations	The number of add operations. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Memory_Remove_Operations	The number of remove operations. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Name	The name of the dynamic memory balancer instance [Sequential = ID Non-Sequential = ID]
Removed_Memory	The amount of memory removed from virtual machines (VM). This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

Hyper-V Statistics

System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Task Manager Detail
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Task Manager Detail
Open Table Name:	HVTASKMGRDETAIL
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = AVG]
Add_Resources_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Add_Resources_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Add_Resources_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Apply_Snapshot_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Apply_Snapshot_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]

Apply_Snapshot_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Clone_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Clone_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Clone_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Create_VSS_Snapshot_Set_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Create_VSS_Snapshot_Set_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Create_VSS_Snapshot_Set_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Define_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Define_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Define_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Destroy_Snapshot_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Destroy_Snapshot_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]

Hyper-V Statistics

Destroy_Snapshot_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Destroy_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Destroy_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Destroy_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Export_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Export_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Export_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Import_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Import_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Import_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = AVG]
Merge_Disk_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]

Merge_Disk_Tasks_in_Progress	<p>The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Merge_Disk_Tasks_Recent_Time	<p>The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Migrate_Virtual_Machine_Tasks_Completed	<p>The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
Migrate_Virtual_Machine_Tasks_in_Progress	<p>The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Migrate_Virtual_Machine_Tasks_Recent_Time	<p>The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Modify_Resources_Virtual_Machine_Tasks_Completed	<p>The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
Modify_Resources_Virtual_Machine_Tasks_in_Progress	<p>The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Modify_Resources_Virtual_Machine_Tasks_Recent_Time	<p>The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Modify_Service_Settings_Tasks_Completed	<p>The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
Modify_Service_Settings_Tasks_in_Progress	<p>The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Modify_Service_Settings_Tasks_Recent_Time	<p>The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Modify_Virtual_Machine_Tasks_Completed	<p>The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>

Hyper-V Statistics

Modify_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Modify_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Pause_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Pause_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Pause_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Remove_Resources_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Remove_Resources_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Remove_Resources_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Reset_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Reset_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Reset_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Restore_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]

Restore_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Restore_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Resume_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Resume_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Resume_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = ID]
Save_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Save_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Save_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Shutdown_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Shutdown_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Shutdown_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]

Hyper-V Statistics

Snapshot_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Snapshot_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Snapshot_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Start_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Start_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Start_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
System	The name of the Hyper-V system. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = ID]
Waiting_to_Start_Virtual_Machine_Tasks_Completed	The number of tasks completed. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Waiting_to_Start_Virtual_Machine_Tasks_in_Progress	The number of tasks that are in progress. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Waiting_to_Start_Virtual_Machine_Tasks_Recent_Time	The time spent to execute the last operation of this type in 100 nanosecond units. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual Machine Bus
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Virtual Machine Bus
Open Table Name:	HVVMBUS
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interrupts_Received	The number of interrupts received. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Interrupts_Received/sec	The number of interrupts received per second [Sequential = DIV Non-Sequential = SUM]
Interrupts_Sent	The number of interrupts sent. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Interrupts_Sent/sec	The number of interrupts sent per second [Sequential = DIV Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Throttle_Events	The number of times that any partition has been throttled. A partition is throttled when its interrupts are disabled. [Sequential = SUM Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual Machine Health Summary
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Virtual Machine Health Summary
Open Table Name:	HVVMHEALTHSUM
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Health_Critical	The number of virtual machines with critical health [Sequential = LST Non-Sequential = SUM]
Health_Ok	The number of virtual machines with OK health [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual Machine Summary
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Virtual Machine Summary
Open Table Name:	HVVMSUM
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = AVG]
Applying_Snapshot	The number of virtual machines that are applying snapshots. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Deleting	The number of virtual machines that are deleting virtual machine information. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Deleting_Saved_State	The number of virtual machines that are deleting saved states. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Deleting_Snapshot	The number of virtual machines that are deleting snapshots. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Exporting	The number of virtual machines that are exporting virtual machine information. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = AVG]

Merging_Disks	<p>The number of virtual machines that are merging disks. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Paused	<p>The number of virtual machines that are paused. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Pausing	<p>The number of virtual machines that are pausing. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Resetting	<p>The number of virtual machines that are resetting. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Resuming	<p>The number of virtual machines that are resuming. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Running	<p>The number of virtual machines that are running. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = ID]</p>
Saved	<p>The number of virtual machines that are saved. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Saving	<p>The number of virtual machines that are saving. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Starting	<p>The number of virtual machines that are starting. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Stopping	<p>The number of virtual machines that are stopping. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
System	<p>The name of the Hyper-V system. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2.</p> <p>[Sequential = ID Non-Sequential = ID]</p>

Taking_Snapshot	The number of virtual machines that are taking snapshots. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Time	The timestamp of the data sample. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = ID]
Total_VMs	The total number of virtual machines. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Turned_Off	The number of virtual machines that are turned off. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]
Waiting_to_Start	The number of virtual machines that are waiting to start. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = LST Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual Network Adapter
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Virtual Network Adapter
Open Table Name:	HVVIRTNETADAPTER
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Broadcast_Packets_Received/sec	The number of broadcast packets received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]

Hyper-V Statistics

Broadcast_Packets_Sent/sec	The number of broadcast packets sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Bytes_Received/sec	The number of bytes received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Bytes_Sent/sec	The number of bytes sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Bytes/sec	The number of bytes sent or received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Directed_Packets_Received/sec	The number of directed packets received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Directed_Packets_Sent/sec	The number of directed packets sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Dropped_Packets_Incoming/sec	The total number of incoming packets dropped per second by the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Dropped_Packets_Outgoing/sec	The total number of outgoing packets dropped per second by the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Extensions_Dropped_Packets_Incoming/sec	The total number of incoming packets dropped per second by switch extensions of the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Extensions_Dropped_Packets_Outgoing/sec	The total number of outgoing packets dropped per second by switch extensions of the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
GUID	The globally unique identifier (GUID) of the Virtual Network Adapter [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
IPsec_offload_Bytes_Receive/sec	The total number of Psec offload bytes received per second by the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
IPsec_offload_Bytes_Sent/sec	The total number of Psec offload bytes sent per second by the network adapter. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]

Multicast_Packets_Received/sec	The number of multicast packets received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Multicast_Packets_Sent/sec	The number of multicast packets sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Name	The name of the virtual network adapter instance [Sequential = ID Non-Sequential = ID]
Packets_Received/sec	The number of packets received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Packets_Sent/sec	The number of packets sent by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Packets/sec	The number of packets sent or received by the network adapter per second [Sequential = DIV Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Hyper-V Statistics

Table Field Hierarchy

Class:	Hyper-V
Subclass:	VM Vid Numa Node
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.VM Vid Numa Node
Open Table Name:	HVVMVIDNUMANODE
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
PageCount	The number of physical pages detected on the Non-Uniform Memory Access (NUMA) node [Sequential = LST Non-Sequential = SUM]
Preferred_NUMA_Node	The preferred Non-Uniform Memory Access (NUMA) node index associated with the partition [Sequential = ID Non-Sequential = ID]
ProcessorCount	The number of logical processors detected on the Non-Uniform Memory Access (NUMA) node [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

5.4. I/O Statistics

The Hyper-V Agent stores I/O statistics in the performance database tables.

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual IDE Controller
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.Virtual IDE Controller
Open Table Name:	HVVIRTIDECONTROLLER
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Read_Bytes/sec	The number of bytes read per second from the disks attached to the Integrated Drive Electronics (IDE) controller [Sequential = DIV Non-Sequential = SUM]
Read_Sectors/sec	The number of sectors read per second from the disks attached to the Integrated Drive Electronics (IDE) controller [Sequential = DIV Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]

Hyper-V Statistics

Write_Bytes/sec	The number of bytes written per second to the disks attached to the Integrated Drive Electronics (IDE) controller [Sequential = DIV Non-Sequential = SUM]
Write_Sectors/sec	The number of sectors written per second to the disks attached to the Integrated Drive Electronics (IDE) controller [Sequential = DIV Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual IDE Controller (Emulated)
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.Virtual IDE Controller (Emulated)
Open Table Name:	HVVIRTIDECONTROLLER
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Read_Bytes/sec	The number of bytes read per second from the disks attached to the Integrated Drive Electronics (IDE) controller [Sequential = DIV Non-Sequential = SUM]
Read_Sectors/sec	The number of sectors read per second from the disks attached to the Integrated Drive Electronics (IDE) controller [Sequential = DIV Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]

Write_Bytes/sec	The number of bytes written per second to the disks attached to the Integrated Drive Electronics (IDE) controller [Sequential = DIV Non-Sequential = SUM]
Write_Sectors/sec	The number of sectors written per second to the disks attached to the Integrated Drive Electronics (IDE) controller [Sequential = DIV Non-Sequential = SUM]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual Storage Device
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Virtual Storage Device
Open Table Name:	HVVIRTSTORDEV
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Error_Count	The number of errors that have occurred on the virtual device [Sequential = SUM Non-Sequential = SUM]
Flush_Count	The number of flush operations that have occurred on the virtual device [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Name	The name of the virtual storage device instance [Sequential = ID Non-Sequential = ID]
Read_Bytes/sec	The number of bytes that have been read per second from the virtual device [Sequential = DIV Non-Sequential = SUM]
Read_Count	The number of read operations that have occurred on the virtual device [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Write_Bytes/sec	The number of bytes written per second to the virtual device [Sequential = DIV Non-Sequential = SUM]
Write_Count	The number of write operations that have occurred to the virtual device [Sequential = SUM Non-Sequential = SUM]

5.5. Network Statistics

The Hyper-V Agent stores network statistics in the performance database tables.

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Legacy Network Adapter
IT Resource Name:	/TeamQuest/System/Hyper-V/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Hyper-V.Legacy Network Adapter
Open Table Name:	HVLEGACYNETADAPTER
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Bytes_Dropped	The number of bytes dropped by the network adapter [Sequential = SUM Non-Sequential = SUM]
Bytes_Received/sec	The number of bytes received per second on the network adapter [Sequential = DIV Non-Sequential = SUM]
Bytes_Sent/sec	The number of bytes sent per second over the network adapter [Sequential = DIV Non-Sequential = SUM]
Frames_Dropped	The number of frames dropped by the network adapter [Sequential = SUM Non-Sequential = SUM]
Frames_Received/sec	The number of frames received per second on the network adapter [Sequential = DIV Non-Sequential = SUM]

Frames_Sent/sec	The number of frames sent per second over the network adapter [Sequential = DIV Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Name	The name of the legacy network adapter instance [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the Hyper-V virtual machine [Sequential = ID Non-Sequential = ID]

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual Switch
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Virtual Switch
Open Table Name:	HVVIRTSWTC
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Broadcast_Packets_Received/sec	The number of broadcast packets received by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Broadcast_Packets_Sent/sec	The number of broadcast packets sent by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Bytes_Received/sec	The number of bytes received by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Bytes_Sent/sec	The number of bytes sent by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]

Hyper-V Statistics

Bytes/sec	The number of bytes sent or received by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Directed_Packets_Received/sec	The number of directed packets received by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Directed_Packets_Sent/sec	The number of directed packets sent by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Dropped_Packets_Incoming/sec	The total number of incoming packets dropped per second by the virtual switch. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Dropped_Packets_Outgoing/sec	The total number of outgoing packets dropped per second by the virtual switch. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Extensions_Dropped_Packets_Incoming/sec	The total number of incoming packets dropped per second by the virtual switch extensions. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Extensions_Dropped_Packets_Outgoing/sec	The total number of outgoing packets dropped per second by the virtual switch extensions. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Learned_Mac_Addresses	The number of learned Media Access Control (MAC) addresses of the virtual switch [Sequential = DIV Non-Sequential = SUM]
Learned_Mac_Addresses/sec	The number of Media Access Control (MAC) addresses learned per second by the virtual switch [Sequential = DIV Non-Sequential = SUM]
Multicast_Packets_Received/sec	The number of multicast packets received by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Multicast_Packets_Sent/sec	The number of multicast packets sent by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Name	The name of the virtual switch instance [Sequential = ID Non-Sequential = ID]
Number_of_Send_Channel_Moves/sec	The total number of send channel moves per second on the virtual switch. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]
Number_of_VMQ_Moves/sec	The total number of virtual machine queue (VMQ) moves per second on the virtual switch. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]

Packets_Flooded	The number of packets flooded by the virtual switch [Sequential = DIV Non-Sequential = SUM]
Packets_Flooded/sec	The number of packets flooded per second by the virtual switch [Sequential = DIV Non-Sequential = SUM]
Packets_Received/sec	The number of packets received by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Packets_Sent/sec	The number of packets sent by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Packets/sec	The number of packets sent or received by the virtual switch per second [Sequential = DIV Non-Sequential = SUM]
Purged_Mac_Addresses	The number of purged Media Access Control (MAC) addresses of the virtual switch [Sequential = LST Non-Sequential = SUM]
Purged_Mac_Addresses/sec	The number of Media Access Control (MAC) addresses purged per second by the virtual switch [Sequential = DIV Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Hyper-V Statistics

Table Field Hierarchy

Class:	Hyper-V
Subclass:	Virtual Switch Port
IT Resource Name:	/TeamQuest/System/Hyper-V/Hosts/ <i>hostname</i>
TeamQuest Table Name:	Hyper-V.Virtual Switch Port
Open Table Name:	HVVIRTSWTCHPORT
Collection interval:	300 seconds (default)
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Broadcast_Packets_Received/sec	The number of broadcast packets received by the virtual switch port per second [Sequential = DIV Non-Sequential = SUM]
Broadcast_Packets_Sent/sec	The number of broadcast packets sent by the virtual switch port per second [Sequential = DIV Non-Sequential = SUM]
Bytes_Received/sec	The number of bytes received by the virtual switch port per second [Sequential = DIV Non-Sequential = SUM]
Bytes_Sent/sec	The number of bytes sent by the virtual switch port per second [Sequential = DIV Non-Sequential = SUM]
Bytes/sec	The number of bytes sent or received by the virtual switch port per second [Sequential = DIV Non-Sequential = SUM]
Directed_Packets_Received/sec	The number of directed packets received by the virtual switch port per second [Sequential = DIV Non-Sequential = SUM]
Directed_Packets_Sent/sec	The number of directed packets sent by the virtual switch port per second [Sequential = DIV Non-Sequential = SUM]
Dropped_Packets_Incoming/sec	The total number of incoming packets dropped per second by the virtual switch port. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = DIV Non-Sequential = SUM]

Dropped_Packets_Outgoing/sec	The total number of outgoing packets dropped per second by the virtual switch port. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = AVG Non-Sequential = SUM]
Extensions_Dropped_Packets_Incoming/sec	The total number of incoming packets dropped per second by the virtual switch extensions by the virtual switch port. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = AVG Non-Sequential = SUM]
Extensions_Dropped_Packets_Outgoing/sec	The total number of outgoing packets dropped per second by the virtual switch extensions by the virtual switch port. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
IPsec_offload_Bytes_Receive/sec	The total number of IPsec offload bytes received per second by the virtual switch port. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = AVG Non-Sequential = SUM]
IPsec_offload_Bytes_Sent/sec	The total number of IPsec offload bytes sent per second by the virtual switch port. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = AVG Non-Sequential = SUM]
IPsec_SAs_Offloaded	The total number of IPsec Security Associations currently offloaded by the virtual switch port. This statistic is not available for Microsoft Windows 2008 or Microsoft Windows 2008 R2. [Sequential = SUM Non-Sequential = SUM]
Multicast_Packets_Received/sec	The number of multicast packets received by the virtual switch port per second [Sequential = AVG Non-Sequential = SUM]
Multicast_Packets_Sent/sec	The number of multicast packets sent by the virtual switch port per second [Sequential = AVG Non-Sequential = SUM]
Name	The name of the virtual switch port instance [Sequential = ID Non-Sequential = ID]
Packets_Received/sec	The number of packets received by the virtual switch port per second [Sequential = AVG Non-Sequential = SUM]
Packets_Sent/sec	The number of packets sent by the virtual switch port per second [Sequential = AVG Non-Sequential = SUM]
Packets/sec	The number of packets sent or received by the virtual switch port per second [Sequential = AVG Non-Sequential = SUM]

Hyper-V Statistics

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the Hyper-V system [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual Switch	The name of the associated virtual switch [Sequential = ID Non-Sequential = ID]

Section 6

IBM AIX Systems

Statistics for IBM AIX systems are collected by the TeamQuest collection agents.

This section contains a listing of the statistics collected for the system:

- System Activity Statistics (see 6.1)
- Disk Space Statistics (see 6.2)
- Network Statistics (see 6.3)
- Workload Manager Statistics (see 6.4)
- Workload Statistics (see 6.5)
- LPAR Configuration Statistics (see 6.6)
- Process Statistics (see 6.7)
- Hardware Inventory Statistics (see 6.8)
- System Log Statistics (see 6.9)
- General Log Statistics (see 6.10)
- TeamQuest Log Statistics (see 6.11)
- Derived Statistics (see 6.12)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV =Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

6.1. System Activity Statistics

The System Activity Agent is used to collect a wide variety of important system statistics. Major resources monitored by this agent include processors, memory, disks, and the operating system kernel.

Note: *The collection of disk input and output (i/o) statistics is disabled by default in current releases of IBM AIX. If the system is not collecting disk i/o history statistics, Block Device statistics are not collected. This includes Block Device.by Device and Block Device.Summary statistics. To enable the collection of these statistics, run the following command:*

```
chdev -l sys0 -a iostat=true
```

Special Processing When Using Sequential Consolidation Method

Special processing occurs when certain records in the Block Device.by Device table are consolidated using the Sequential consolidation method. The following formulas are used to calculate the %busy, Actual_Interval, and record_count statistic values:

%busy

The %busy field uses a new consolidation method that uses the following formula to produce the consolidated %busy value:

$$\text{consolidated \%busy} = \%busy * \text{record_count} * \text{Actual_Interval}$$

At the end of the aggregation processing step after multiple records have been combined together to produce a single consolidated record, the %busy field contains the consolidated %busy value.

An additional processing step is performed using the following formula to produce a final %busy value that is stored into the consolidated record:

$$\%busy = \frac{\text{consolidated \%busy}}{\text{record_count} * \text{Interval}}$$

Note: The record_count field value used in the above formula must have already been generated using the record_count formula.

Block Device.by Device table records that have been stored by previous levels of TeamQuest collection agents do not contain the record_count field. For these records, a value of 1 is assumed for the record_count value.

Actual_Interval

For consolidated records (both reduced and not reduced), the Actual_Interval field should contain the Interval value at the end of the aggregation processing step.

record_count

The record_count field value is updated at the end of the aggregation processing step using the following formula:

$$\text{record_count} = \frac{(\sum(\text{record_count} * \text{Actual_Interval})) + (\text{Interval} - \sum \text{Actual_Interval})}{\text{Interval}}$$

The first part of the numerator is a summation of the record_count value multiplied by the Actual_Interval value across all of the consolidated records. This value should exist in the record_count field at the end of the aggregation processing step since the record_count field is weighted by Actual_Interval.

The Actual_Interval value used in the second part of the numerator is a summation of all the Actual_Interval values of the consolidated records. This value should exist in the Actual_Interval field at the end of the aggregation processing step since Actual_Interval is a summation.

The Interval value in the formula is the time range in seconds that the final consolidated record represents. For example, if 5-minute records are generated, the Interval value is calculated as 5 multiplied by 60.

Parameter Hierarchy

Class:	Block Device
Subclass:	by Device
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Block Device.by Device
Open Table Name:	BLKDEVBYDEVICE
Resource:	disk0, disk1, ...
Statistic Name:	
%busy	The percentage of time this device was servicing a transfer request [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-util.rpt
actq_avwait*	The average run queue wait time in milliseconds
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgresp*	The average response time of an I/O on a device. Calculated as avwait + avserv
avque	The average number of requests outstanding [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/io/dsk-q.rpt

avserv	<p>The average time in milliseconds to service each transfer request (includes seek, rotation latency, and data transfer times) for the device</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/io/dsk-time.rpt</p>
avwait	<p>The average time in milliseconds that transfer requests are idle in the queue while the queue is occupied</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/io/dsk-time.rpt</p>
Interval	<p>The expected sampling interval in seconds</p> <p>[Sequential = SUM Non-Sequential = ID]</p>
IO_intensity*	<p>The activity of an I/O device. This is the product of the I/O response time in milliseconds and the I/O transfer rate in I/Os per second. This is proportional to the average queue length (the number of I/O requests waiting or in progress at the I/O device).</p>
Kbytes/s	<p>The rate at which data is transferred in kilobytes per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/hp-ux/sys-act/io/dsk-xfer.rpt</p> <p>/report/hp-ux/sys-act/io/top-dsk.rpt</p>
record_count	<p>The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
reduction_name	<p>The name of reduction rule</p> <p>[Sequential = ID Non-Sequential = ID]</p>
reduction_source	<p>The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
transfers/s	<p>The number of physical transfers to and from the disk per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/hp-ux/sys-act/io/dsk-xfer.rpt</p>
waitq_avwait*	<p>The average wait queue wait time in milliseconds</p>

Class:	Block Device
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Block Device.Summary
Open Table Name:	BLKDEVSUM
Statistic Name:	
transfers/s	The number of physical transfers to and from all devices per second [Sequential = AVG Non-Sequential = SUM]
Class:	CPU
Subclass:	by Processor
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.by Processor
Open Table Name:	CPUBYPROC
Resource:	cpu0, cpu1, ...
Statistic Name:	
%idle	The percentage of CPU time spent idle for this CPU [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/cpu/per-cpu.rpt
%sys	The percentage of CPU time spent running in system mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/cpu/per-cpu.rpt
%usr	The percentage of CPU time spent running in user mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/cpu/per-cpu.rpt
%wio	The percentage of CPU time spent idle while some process is waiting for I/O completion [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/cpu/per-cpu.rpt
pct_fardisp	The percentage of thread dispatches to this processor that were far dispatches. This statistic is not available on systems before POWER7 architectures. [Sequential = AVG Non-Sequential = AVG] Note: The hardware meanings for local, near, and far vary with different architectures.

pct_localdisp	<p>The percentage of thread dispatches to this processor that were local dispatches. This statistic is not available on systems before POWER7 architectures.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>Note: <i>The hardware meanings for local, near, and far vary with different architectures.</i></p>
pct_neardisp	<p>Percentage of thread dispatches to this processor that were near dispatches. This statistic is not available on systems before POWER7 architectures.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>Note: <i>The hardware meanings for local, near, and far vary with different architectures.</i></p>

Notes:

- *The statistics reported in the CPU.by LPAR table on a shared logical partition are relative to either the entitlement of the partition or to the actual physical processor usage of the partition, whichever is greater.*
- *The statistics reported in the CPU.by LPAR table on a dedicated logical partition are relative to the configured capacity of the partition.*
- *For more information on partition entitlement, shared partitions, or dedicated partitions, see the IBM PowerVM Virtualization documentation.*

Class:	CPU
Subclass:	by LPAR
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	CPU.by LPAR
Open Table Name:	CPUBYLPAR
Statistic Name:	
%entc	<p>The percentage of the entitled processor capacity consumed. Available only in a shared dynamic logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
%idle	<p>The percentage of the entitled processing capacity unused while the partition was idle and did not have any outstanding disk I/O requests. Available only in a logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/aix/sys-act/cpu/lpar-cpu.rpt</p>
%lpar_phys_busy	<p>The percentage of the processor capacity consumed. If an LPAR moves into or out of the shared adapter pool, this statistic may be displayed as <N/A>. Available only in a shared dynamic logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>

%lpar_pool_busy	<p>The percentage of the physical processor pool capacity consumed. If an LPAR moves into or out of the shared adapter pool, this statistic may be displayed as <N/A>. Available only in a shared dynamic logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
%sys	<p>The percentage of the entitled processing capacity used while executing at the system level. Available only in a logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/aix/sys-act/cpu/lpar-cpu.rpt</p>
%total_pool_busy	<p>The percentage of the processor pool capacity consumed. Allow performance information collection must be enabled. If the Allow performance information collection is not enabled, this statistic is displayed as <N/A>. If an LPAR moves into or out of the shared adapter pool, this statistic may be displayed as <N/A>. Available only in a shared dynamic logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>Note: In earlier versions of hardware management console (HMC), the Allow performance information collection name is Allow shared processor pool utilization.</p>
%user	<p>The percentage of the entitled processing capacity used while executing at the user level. Available only in a logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/aix/sys-act/cpu/lpar-cpu.rpt</p>
%wait	<p>The percentage of the entitled processing capacity unused while the partition was idle and had outstanding disk I/O requests. Available only in a logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/aix/sys-act/cpu/lpar-cpu.rpt</p>
app	<p>The number of available processors in the shared pool. Allow performance information collection must be enabled. If the Allow performance information collection is not enabled, this statistic is displayed as <N/A>. If an LPAR moves into or out of the shared adapter pool, this statistic may be displayed as <N/A>. Available only in a shared dynamic logical partition environment.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>Note: In earlier versions of hardware management console (HMC), the Allow performance information collection name is Allow shared processor pool utilization.</p>
ivcsw	<p>The number of involuntary virtual context switches per second. An involuntary context switch occurs when a process is interrupted by a higher priority process or exceeds its time allotment.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>

pct_fardisp	<p>The percentage of thread dispatches that were far dispatches. This statistic is not available on systems before POWER7 architectures. [Sequential = AVG Non-Sequential = AVG]</p> <p>Note: <i>The hardware meanings for local, near, and far vary with different architectures.</i></p>
pct_localdisp	<p>The percentage of thread dispatches that were local dispatches. This statistic is not available on systems before POWER7 architectures. [Sequential = AVG Non-Sequential = AVG]</p> <p>Note: <i>The hardware meanings for local, near, and far vary with different architectures.</i></p>
pct_neardisp	<p>Percentage of thread dispatches that were near dispatches. This statistic is not available on systems before POWER7 architectures. [Sequential = AVG Non-Sequential = AVG]</p> <p>Note: <i>The hardware meanings for local, near, and far vary with different architectures.</i></p>
phint	<p>The number of phantom (targeted to another shared partition in this pool) interruptions received. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = SUM]</p>
physc	<p>The number of physical processors consumed. Available only in a logical partition environment. [Sequential = AVG Non-Sequential = AVG]</p>
vcsww	<p>The number of virtual context switches that are virtual processor hardware preemptions. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = SUM]</p>
vvcsww	<p>The number of voluntary virtual context switches per second. A voluntary context switch usually occurs when the process has nothing to do or when a process is waiting for an event to happen (for example, an I/O operation to complete). [Sequential = AVG Non-Sequential = SUM]</p>

The CPU.RelativePerformance table is populated for physical systems only. It is not populated for AIX logical partitions (LPARs), VMware guests, Hyper-V guests, Solaris logical domains (LDOMs), Solaris guest LDOMs, KVM guests, and Linux on POWER systems.

Table Field Hierarchy

Class:	CPU
Subclass:	RelativePerformance
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
rel_unused	The amount of CPU resources not used based on a common, relative scale. This statistic is not available for IBM AIX systems. [Sequential = AVG Non-Sequential = SUM]
rel_used	The amount of CPU resources used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Note: *The statistics reported in the CPU.Summary table are relative to the usage of online logical processors in the logical partition. For more information on logical processors, see the IBM PowerVM Virtualization documentation.*

Class:	CPU
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/systemname/CPU
TeamQuest Table Name:	CPU.Summary
Open Table Name:	CPUSUM
Statistic Name:	
%busy	The percentage of time the CPU was not idle [Sequential = AVG Non-Sequential = AVG]
%idle	The percentage of total CPU time spent idle while no processes are waiting for I/O completion [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/cpu/cpu-util.rpt
%sys	The percentage of total CPU time spent running in system mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/cpu/cpu-util.rpt
%usr	The percentage of total CPU time spent running in user mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/cpu/cpu-util.rpt
%wio	The percentage of total CPU time spent idle while some processes are waiting for I/O completion [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/cpu/cpu-util.rpt
online_cpus	The number of CPUs that were online at the end of the sampling interval [Sequential = LST Non-Sequential = SUM]

Note: The IBM POWER.Block Device Summary table is a derived table and therefore is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	IBM POWER
Subclass:	Block Device Summary
IT Resource Name:	/TeamQuest/System/systemname
Table type:	Performance

Statistic Name	Description
Frame	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
mpid	The identifier of the memory pool that the LPAR belongs to [Non-Sequential = ID]
partition_group_id	The LPAR group that the LPAR is a member of [Non-Sequential = ID]
partition_name	The logical partition name as assigned by the management entity [Sequential = LST Non-Sequential = ID]
shared_pool_id	The identifier of the shared pool of physical processors that the LPAR is a member of [Non-Sequential = ID]
transfers/s	The number of physical transfers to and from all devices per second [Sequential = AVG Non-Sequential = SUM]

Note: The IBM POWER.CPU by LPAR table is a derived table and therefore is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	IBM POWER
Subclass:	CPU by LPAR
IT Resource Name:	/TeamQuest/System/systemname
Table type:	Performance

Statistic Name	Description
%entc	The percentage of the entitled processor capacity consumed. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = AVG]
%idle	The percentage of the entitled processing capacity unused while the partition was idle and did not have any outstanding disk I/O requests. Available only in a logical partition environment. [Sequential = AVG Non-Sequential = AVG]

<code>%lpar_phys_busy</code>	The percentage of the processor capacity consumed. If an LPAR moves into or out of the shared adapter pool, this statistic may be displayed as <N/A>. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = SUM]
<code>%lpar_pool_busy</code>	The percentage of the physical processor pool capacity consumed. If an LPAR moves into or out of the shared adapter pool, this statistic may be displayed as <N/A>. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = SUM]
<code>%sys</code>	The percentage of the entitled processing capacity used while executing at the system level. Available only in a logical partition environment. [Sequential = AVG Non-Sequential = AVG]
<code>%total_pool_busy</code>	The percentage of the processor pool capacity consumed. Allow shared processor pool utilization must be enabled. If an LPAR moves into or out of the shared adapter pool, this statistic may be displayed as <N/A>. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = SUM]
<code>%user</code>	The percentage of the entitled processing capacity used while executing in user mode. Available only in a logical partition environment. [Sequential = AVG Non-Sequential = AVG]
<code>%wait</code>	The percentage of the entitled processing capacity unused while the partition was idle and had outstanding disk I/O requests. Available only in a logical partition environment. [Sequential = AVG Non-Sequential = AVG]
<code>app</code>	The number of available processors in the shared pool. Allow shared processor pool utilization must be enabled. If an LPAR moves into or out of the shared adapter pool, this statistic may be displayed as <N/A>. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = SUM]
<code>Frame</code>	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
<code>mpid</code>	The identifier of the memory pool that the LPAR belongs to [Non-Sequential = ID]
<code>partition_group_id</code>	The LPAR group that the LPAR is a member of [Non-Sequential = ID]
<code>partition_name</code>	The logical partition name as assigned by the management entity [Sequential = LST Non-Sequential = ID]
<code>phint</code>	The number of phantom (targeted to another shared partition in this pool) interrupts received. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = SUM]

physc	The number of physical processors consumed. Available only in a logical partition environment. [Sequential = AVG Non-Sequential = AVG]
shared_pool_id	The identifier of the shared pool of physical processors that the LPAR is a member of [Non-Sequential = ID]
vcsww	The number of virtual context switches that are virtual processor hardware preemptions. Available only in a shared dynamic logical partition environment. [Sequential = AVG Non-Sequential = SUM]
vvcsww	The number of voluntary virtual context switches per second. A voluntary context switch usually occurs when the process has nothing to do or when a process is waiting for an event to happen (for example, when an I/O operation completes). [Sequential = AVG Non-Sequential = SUM]

Note: The IBM POWER.CPU Relative Performance table is a derived table and therefore is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	IBM POWER
Subclass:	CPU Relative Performance
IT Resource Name:	/TeamQuest/System/systemname
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Sequential = AVG Non-Sequential = SUM]
Frame	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
mpid	The identifier of the memory pool that the logical partition (LPAR) belongs to [Non-Sequential = ID]
partition_group_id	The identifier of the LPAR group that the LPAR is a member of [Non-Sequential = ID]
partition_name	The logical partition name as assigned by the management entity [Sequential = LST Non-Sequential = ID]

rel_unused	<p>The amount of CPU resources not used based on a common, relative scale. CPU available to an LPAR can vary over time due to configuration (shared versus dedicated CPU for the LPAR), donation of dedicated CPU by the LPAR, usage of shared CPU pools, and so on. Unused relative CPU is not calculated and is reported as <N/A>. This statistic is not available for IBM AIX systems.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
rel_used	<p>The amount of CPU resources used based on a common, relative scale</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
shared_pool_id	<p>The identifier of the shared pool of physical processors that the LPAR is a member of</p> <p>[Non-Sequential = ID]</p>
System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>

The PowerVM CEC Agent retrieves configuration information for the logical partitions hosted on an IBM POWER server. The information is stored at startup and updated once a day in the IBM POWER.FRAME CONFIG table. In addition, the PowerVM CEC Agent checks for partition configuration changes at a specified interval and records the configuration if changes are detected. This interval is controlled by the Performance Data Policy applied to the database and defaults to 1-hour.

Table Field Hierarchy

Class:	IBM POWER
Subclass:	FRAME CONFIG
IT Resource Name:	/TeamQuest/System/ <i>framename</i>
TeamQuest Table Name:	IBM POWER.FRAME CONFIG
Open Table Name:	IBMPOWERFRAMECONFIG
Collection interval:	N/A
Default retentions:	1 year
Table type:	Performance

Statistic Name	Description
Frame	The name of the physical system as displayed by the management entity [Sequential = ID Non-Sequential = ID]
managing_node	The name of the entity from which the configuration data was retrieved as given by the PowerVM CEC Policy that is applied to the node. This name is typically a management entity or a VIOS that hosts an Integrated Virtualization Manager (IVM). [Sequential = ID Non-Sequential = ID]
partition_name	The name of the partition profile as it was activated. This field is limited to 51 characters. Any name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
partition_number	The number of the logical partition as assigned by the management entity [Sequential = ID Non-Sequential = ID]
partition_state	The current state of the partition profile. For example, Not Activated, Starting, Running, Shutting Down, Error, Open Firmware, or Not Available. [Sequential = ID Non-Sequential = ID]
System	The domain name system (DNS) host name. This field is determined using the Internet Protocol (IP) address used by the management entity for the partition. If the IP address is available, but cannot be mapped to a DNS name, this value is the IP address of the partition. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Note: The IBM POWER.Memory table is a derived table and therefore is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	IBM POWER
Subclass:	Memory
IT Resource Name:	/TeamQuest/System/systemname
Table type:	Performance

Statistic Name	Description
Frame	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
freereal (MB)	The amount of real free memory available in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
mpid	The identifier of the memory pool that the LPAR belongs to [Non-Sequential = ID]
partition_group_id	The LPAR group that the LPAR is a member of [Non-Sequential = ID]
partition_name	The logical partition name as assigned by the management entity [Sequential = LST Non-Sequential = ID]
shared_pool_id	The identifier of the shared pool of physical processors that the LPAR is a member of [Non-Sequential = ID]
totalreal (MB)	The total amount of memory in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
totalvirt (MB)	The total amount of virtual memory in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]

Note: The IBM POWER.Memory Page Space table is a derived table and therefore is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	IBM POWER
Subclass:	Memory Page Space
IT Resource Name:	/TeamQuest/System/systemname
Table type:	Performance

Statistic Name	Description
Frame	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
free (MB)	The amount of free page space in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
mpid	The identifier of the memory pool that the LPAR belongs to [Non-Sequential = ID]
partition_group_id	The LPAR group that the LPAR is a member of [Non-Sequential = ID]
partition_name	The logical partition name as assigned by the management entity [Sequential = LST Non-Sequential = ID]
reserved (MB)	The amount of reserved page space in megabytes [Sequential = AVG Non-Sequential = AVG]
shared_pool_id	The identifier of the shared pool of physical processors that the LPAR is a member of [Non-Sequential = ID]
total (MB)	The total amount of page space in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]

Note: The IBM POWER.Memory Real table is a derived table and therefore is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	IBM POWER
Subclass:	Memory Real
IT Resource Name:	/TeamQuest/System/systemname
Table type:	Performance

Statistic Name	Description
client (MB)	The amount of memory currently allocated to cache remotely mounted files in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
comp (MB)	The amount of real memory currently allocated to computational page frames in megabytes (MB). Computational page frames are generally those that are backed by paging space. [Sequential = AVG Non-Sequential = AVG]
Frame	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
free (MB)	The amount of real free memory available in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
inuse (MB)	The total amount of real memory that is in use in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
mpid	The identifier of the memory pool that the LPAR belongs to [Non-Sequential = ID]
noncomp (MB)	The amount of real memory currently allocated to noncomputational page frames in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
partition_group_id	The LPAR group that the LPAR is a member of [Non-Sequential = ID]
partition_name	The logical partition name as assigned by the management entity [Sequential = LST Non-Sequential = ID]
pinned (MB)	The total amount of real memory that is pinned in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
shared_pool_id	The identifier of the shared pool of physical processors that the LPAR is a member of [Non-Sequential = ID]
total (MB)	The total amount of real memory in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]

Class:	Kernel
Subclass:	Buffers
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Buffers
Open Table Name:	KNLBUFFS
Statistic Name:	
%rcache	The percentage of logical reads satisfied from the buffer cache [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/kernel/bufc-hit.rpt
%wcache	The percentage of logical writes satisfied from the buffer cache [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/kernel/bufc-hit.rpt
bread/s	The number of reads per second between system buffers and block devices [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/bufc-xfr.rpt
bwrit/s	The number of writes per second between system buffers and block devices [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/bufc-xfr.rpt
lread/s	The number of read accesses of system buffers (logical reads) per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/bufc-xfr.rpt
lwrit/s	The number of write accesses of system buffers (logical writes) per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/bufc-xfr.rpt
pread/s	The number of physical read requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/phys-xfr.rpt
pwrit/s	The number of physical write requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/phys-xfr.rpt

Class:	Kernel
Subclass:	File Access
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.File Access
Open Table Name:	KNLFILEACCESS
Statistic Name:	
dirblk/s	The number of directory block reads issued per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/f-access.rpt
iget/s	The number of files located by i-node entries per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/f-access.rpt
namei/s	The number of file system path searches per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/f-access.rpt
Class:	Kernel
Subclass:	IPC (inter process communication)
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.IPC
Open Table Name:	KNLIPC
Statistic Name:	
msg/s	The number of message operations (sends and receives) per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/msg-sema.rpt
sema/s	The number of semaphore operations per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/msg-sema.rpt

Class:	Kernel
Subclass:	Load Average
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernal.Load Average
Open Table Name:	KNLLOADAVG
Statistic Name:	
1 min	The average number of runnable processes in the last one-minute interval that were waiting for system resources [Sequential = LST Non-Sequential = SUM]
5 min	The average number of runnable processes in the last five-minute interval that were waiting for system resources [Sequential = LST Non-Sequential = SUM]
15 min	The average number of runnable processes in the last fifteen-minute interval that were waiting for system resources [Sequential = LST Non-Sequential = SUM]
Class:	Kernel
Subclass:	Paging
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Kernel.Paging
Open Table Name:	KNLPAGING
Statistic Name:	
backtracks/s	The number of page faults per second that occurred while resolving previous page faults [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/pageflt.rpt
cycles/s	The number of clock hand cycles per second [Sequential = AVG Non-Sequential = SUM]
execfills/s	The number of instruction page faults satisfied per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/pageflt.rpt
freewaits/s	The number of processes per second waiting for free frames to be gathered [Sequential = AVG Non-Sequential = SUM]
iodones/s	The number of virtual memory manager (VMM) read and write I/O requests completed per second [Sequential = AVG Non-Sequential = SUM]
iostarts/s	The number of virtual memory manager (VMM) read and write I/O requests started per second [Sequential = AVG Non-Sequential = SUM]
lockmisses/s	The number of page faults caused by lock misses per second [Sequential = AVG Non-Sequential = SUM]

pendwaits/s	The number of processes waiting for a page-in I/O to complete per second [Sequential = AVG Non-Sequential = SUM]
pflts/s	The number of page faults per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/pageflt.rpt
pgin/s	The number of pages paged in from paging space and file space per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/pagein.rpt
pgout/s	The number of pages paged out to paging space and file space per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/pageout.rpt
pgspin/s	The number of pages paged in from paging space per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/pagein.rpt
pgspout/s	The number of pages paged out to paging space per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/pageout.rpt
rclms/s	The number of page faults satisfied per second without having to initiate new I/O requests [Sequential = AVG Non-Sequential = SUM]
scans/s	The number of pages examined by the clock hand per second [Sequential = AVG Non-Sequential = SUM]
steals/s	The number of page steals per second [Sequential = AVG Non-Sequential = SUM]
zerofills/s	The number of page faults satisfied by zero-filling per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/pageflt.rpt

Class:	Kernel
Subclass:	Queues
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Queues
Open Table Name:	KNLQS
Statistic Name:	
%runocc	The percentage of time the run queue is occupied [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/kernel/q-util.rpt
avg_cpuq_sz	The average length of the run queue per CPU (a queue of processes in memory and runnable) [Sequential = AVG Non-Sequential = AVG]
avg_runq_sz	The average length of the run queue (a queue of processes in memory and runnable) [Sequential = AVG Non-Sequential = AVG] View Reports: /report/aix/sys-act/kernel/q-sizes.rpt /report/aix/sys-act/kernel/runq.rpt
cpuq_sz	The average length of the run queue per CPU (a queue of processes in memory and runnable) while the run queue is occupied [Sequential = AVG Non-Sequential = AVG]
blocked	The average number of threads blocked on I/O (waiting for an I/O to finish) [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/blocked
pswch/s	The number of process switches per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/p-switch.rpt
runq_sz	The average length of the run queue (a queue of processes in memory and runnable) while the run queue is occupied [Sequential = AVG Non-Sequential = AVG] View Reports: /report/aix/sys-act/kernel/q-sizes.rpt /report/aix/sys-act/kernel/runq.rpt

Class:	Kernel
Subclass:	Tables
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Tables
Open Table Name:	KNLTABS
Statistic Name:	
file_sz	The number of entries currently being used in the file table. This count is taken at the end of the sampling interval. [Sequential = AVG Non-Sequential = AVG]
proc_sz	The number of entries currently being used in the process table. This count is taken at the end of the sampling interval. [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/kernel/tbl-size.rpt
thread_sz	The number of entries currently being used in the thread table. This count is taken at the end of the sampling interval. [Sequential = AVG Non-Sequential = AVG]
Class:	Kernel
Subclass:	TTY
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.TTY
Open Table Name:	KNLTTY
Statistic Name:	
canch/s	The number of input characters per second processed by canon (canonical queue) [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/tty-xfer.rpt
mdmin/s	The number of modem interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/tty-intr.rpt
outch/s	The number of output characters transferred per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/tty-xfer.rpt
rawch/s	The number of input characters per second transferred in raw mode [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/tty-xfer.rpt

rcvin/s	The number of receiver hardware interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/tty-intr.rpt
xmtin/s	The number of transmitter hardware interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/kernel/tty-intr.rpt
Class:	Memory
Subclass:	AME
IT Resource Name:	/TeamQuest/System/systemname/Memory
TeamQuest Table Name:	Memory.AME
Open Table Name:	MEMAME
Statistic Name:	
cfr	The amount of free memory in gigabytes in the compressed pool [Sequential = LST Non-Sequential = SUM]
ci/s	The number of decompression operations or page-ins per second from the compressed pool [Sequential = AVG Non-Sequential = SUM]
cmsz	The target size of the compressed memory pool in gigabytes [Sequential = LST Non-Sequential = SUM]
co/s	The number of compression operations or page-outs per second to the compressed pool [Sequential = AVG Non-Sequential = SUM]
cpg	The number of pages in the compressed memory pool [Sequential = LST Non-Sequential = SUM]
cpsz	The compressed size of the pages in gigabytes in the compressed memory pool [Sequential = LST Non-Sequential = SUM]
csz	The size of the compressed memory pool in gigabytes [Sequential = LST Non-Sequential = SUM]
ctsz	The target size of the compressed memory pool in gigabytes [Sequential = LST Non-Sequential = SUM]
cwpg	The number of working storage pages in the compressed pool [Sequential = LST Non-Sequential = SUM]
cwsz	The compressed size of the working pages in gigabytes in the compressed memory pool [Sequential = LST Non-Sequential = SUM]
cx	The achieved memory expansion factor. This is the factor that the system has been able to expand the true memory by. [Sequential = LST Non-Sequential = AVG]
dxm	The size of the expanded memory deficit in gigabytes [Sequential = LST Non-Sequential = SUM]

mem	The expanded memory size in gigabytes of the LPAR [Sequential = LST Non-Sequential = SUM]
tmem	The true memory size in gigabytes of the LPAR [Sequential = LST Non-Sequential = SUM]
txf	The target memory expansion factor. This is expressed as a multiplier that the system attempts to expand the true memory by. A target expansion factor of 2.0 indicates that the target size of the expanded memory is twice the size of the true memory. [Sequential = LST Non-Sequential = AVG]
usz	The size of the uncompressed memory pool in gigabytes [Sequential = LST Non-Sequential = SUM]
Class:	Memory
Subclass:	AMS
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Memory.AMS
Open Table Name:	MEMAMS
Statistic Name:	
hpi/s	The number of Hypervisor page-ins per second for the partition (system). A Hypervisor page-in occurs if a page is being referenced but is not available in real memory because it has been previously paged out by the Hypervisor. [Sequential = AVG Non-Sequential = SUM]
hpit	The average time spent per Hypervisor page-ins in milliseconds [Sequential = AVG Non-Sequential = AVG]
iome	The amount of memory in gigabytes permanently reserved in a memory pool to handle I/O activity [Sequential = LST Non-Sequential = SUM]
loan	The amount of logical memory in gigabytes that is loaned to the Hypervisor [Sequential = LST Non-Sequential = SUM]
mpsz	The size of the shared memory pool in gigabytes [Sequential = LST Non-Sequential = SUM]
memwght	The memory weight. This is a relative number used by the Hypervisor to prioritize the physical memory assignment from the shared memory pool to the logical partition. A higher value increases the probability that more physical memory is assigned to the logical partition. [Sequential = LST Non-Sequential = ID]
pmem	The size of the memory pool backing the logical memory in gigabytes [Sequential = LST Non-Sequential = SUM]

Class: Memory
Subclass: N/A
IT Resource Name: /TeamQuest/System/*systemname*/Memory
TeamQuest Table Name: Memory
Open Table Name: MEM
Statistic Name:
 freereal (MB) The amount of real free memory available in megabytes (MB)
 [Sequential = AVG Non-Sequential = AVG]
 totalreal (MB) The total amount of memory in megabytes (MB)
 [Sequential = AVG Non-Sequential = AVG]
 totalvirt (MB) The total amount of virtual memory in megabytes (MB)
 [Sequential = AVG Non-Sequential = AVG]

Class: Memory
Subclass: Page Space
IT Resource Name: /TeamQuest/System/*systemname*/Memory
TeamQuest Table Name: Memory.Page Space
Open Table Name: MEMPGSPACE
Statistic Name:
 free (MB) The amount of free page space in megabytes (MB)
 [Sequential = AVG Non-Sequential = AVG]
 View Report:
 /report/aix/sys-act/memory/pgspmem.rpt
 reserved (MB) The amount of reserved page space in megabytes (MB)
 [Sequential = AVG Non-Sequential = AVG]
 total (MB) The total amount of page space in megabytes (MB)
 [Sequential = AVG Non-Sequential = AVG]
 View Report:
 /report/aix/sys-act/memory/pgspmem.rpt

Class:	Memory
Subclass:	Real
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Memory.Real
Open Table Name:	MEMREAL
Statistic Name:	
client (MB)	The amount of memory currently allocated to cache remotely mounted files in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
comp (MB)	The amount of real memory currently allocated to computational page frames in megabytes (MB). Computational page frames are generally those that are backed by paging space. [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/memory/memuse.rpt
free (MB)	The amount of real free memory available in megabytes (MB) [Sequential = AVG Non-Sequential = AVG] View Reports: /report/aix/sys-act/memory/realmem.rpt /report/aix/sys-act/memory/memuse.rpt
inuse (MB)	The total amount of real memory that is in use in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
noncomp (MB)	The amount of real memory currently allocated to noncomputational page frames in megabytes (MB) [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/memory/memuse.rpt
pinned (MB)	The total amount of real memory that is pinned in megabytes (MB) [Sequential = AVG Non-Sequential = AVG]
total (MB)	The total amount of real memory in megabytes (MB) [Sequential = AVG Non-Sequential = AVG] View Report: /report/aix/sys-act/memory/realmem.rpt

Class:	System Call
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	System Call.Summary
Open Table Name:	SYSCALLSUM
Statistic Name:	
exec/s	The number of exec system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/syscall/imp-scal.rpt
fork/s	The number of fork system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/syscall/imp-scal.rpt
rchar/s	The number of characters transferred by read system calls in the interval in bytes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/syscall/scal-xfr.rpt
scall/s	The total number of system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/syscall/imp-scal.rpt
sread/s	The number of read system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/syscall/imp-scal.rpt
swrit/s	The number of write system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/syscall/imp-scal.rpt
wchar/s	The number of characters transferred by write system calls in the interval in bytes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/sys-act/syscall/scal-xfr.rpt

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
bsp interval	The number of seconds elapsed between two data samples of the System Activity Agent [Sequential = SUM Non-Sequential = ID]
tqbsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqbsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

Note: *WPAR statistics are only collected when TeamQuest Manager is running under a privileged user ID.*

Class:	WPAR
Subclass:	by WPAR
IT Resource Name:	/TeamQuest/System/systemname/Wpar
TeamQuest Table Name:	WPAR.by WPAR
Open Table Name:	WPARBYWPAR
Statistic Name:	
%sys	The percentage of used CPU time spent running in system mode. This is a percentage of cycles actually used, not a percentage of the available cycles. [Sequential = AVG Non-Sequential = AVG]
%usr	The percentage of used CPU time spent running in user mode. This is a percentage of cycles actually used, not a percentage of the available cycles. [Sequential = AVG Non-Sequential = AVG]
physc	The number of physical processors consumed by the WPAR [Sequential = AVG Non-Sequential = AVG]
Wentc	The percentage of the WPAR's entitled CPU cycles actually consumed [Sequential = AVG Non-Sequential = AVG]

6.2. Disk Space Statistics

Disk space statistics are maintained in the TeamQuest performance database by the Disk Space Agent. The statistics are classified by the hierarchy of key names. A statistic marked with an asterisk (*) is a derived statistic.

Parameter Hierarchy

Class:	Disk Space
Subclass:	by File System
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Disk Space.by File System
Open Table Name:	DISKSPACEBYFILESYS
Resource:	file system1, file system2, ...
Statistic Name:	
%inodes free*	The percentage of i-nodes available (not in use) on the file system at the end of the interval View Report: /report/aix/dskspace/total/pct-inod.rpt
%inodes used*	The percentage of i-nodes in use on the file system at the end of the interval View Report: /report/aix/dskspace/total/pct-inod.rpt
%space free*	The percentage of total space available (not in use) on the file system at the end of the interval View Reports: /report/aix/dskspace/total/pctspace.rpt /report/aix/dskspace/total/low-ones.rpt
%space used*	The percentage of total space in use on the file system at the end of the interval View Reports: /report/aix/dskspace/total/pctspace.rpt /report/aix/dskspace/total/fullest.rpt
%user space free*	The percentage of total user space available (not in use) on the file system at the end of the interval View Reports: /report/aix/dskspace/user/pctspace.rpt /report/aix/dskspace/user/low-ones.rpt
%user space used*	The percentage of total user space in use on the file system at the end of the interval View Reports: /report/aix/dskspace/user/pctspace.rpt /report/aix/dskspace/user/fullest.rpt
capacity	The percentage of total space in use on the file system at the end of the interval [Sequential = LST Non-Sequential = AVG]

free (Mb)	<p>The amount of space available (not in use) on the file system in megabytes (Mb) at the end of the interval including the space held back from normal users</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/aix/dskspace/total/dskspace.rpt</p> <p>/report/aix/dskspace/total/low-ones.rpt</p>
free inodes	<p>The number of available (not in use) i-nodes on the file system at the end of the interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/dskspace/total/i-nodes.rpt</p>
total (Mb)	<p>The total (used + available) amount of space on the file system in megabytes (Mb) at the end of the interval including the space held back from normal users</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/dskspace/total/dskspace.rpt</p>
total inodes	<p>The total (used + available) number of i-nodes on the file system at the end of the interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/dskspace/total/i-nodes.rpt</p>
user free (Mb)	<p>The amount of space available (not in use) on the file system in megabytes (Mb) at the end of the interval not including the space held back from normal users</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/aix/dskspace/user/dskspace.rpt</p> <p>/report/aix/dskspace/user/low-ones.rpt</p>
user total (Mb)*	<p>The total (used + available) amount of space on the file system in megabytes (Mb) at the end of the interval not including the space held back from normal users</p> <p>View Report:</p> <p>/report/aix/dskspace/user/dskspace.rpt</p>

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
dsp interval	The number of seconds elapsed between two data samples of the Disk Space Agent [Sequential = SUM Non-Sequential = ID]
tqdsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqdsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

6.3. Network Statistics

The Network Agent collects information on the interaction of the system with the network. The types of data collected by the agent include network interfaces, NFS and Remote Protocol Call (RPC), and Transmission Control Protocol (TCP).

In NFS environments, the term *server* refers to the system that owns a file system and allows other systems on the network to access the files on the file system. A *client* is a system that mounts a nonlocal file system such as NFS and accesses files from the NFS mounted file system. Any given system could be a server to some file systems and a client to other file systems. It is not uncommon to have systems that export several file systems for other systems to use and mount several other network file systems. All of the NFS statistics collected by this agent show the activity of NFS clients and servers that exist on the system where the agent is run.

Note: *The IBM POWER.Network Summary table is a derived table and therefore is only available for viewing in TeamQuest Analyzer.*

Table Field Hierarchy

Class:	IBM POWER
Subclass:	Network Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
Table type:	Performance

Statistic Name	Description
errors/s	The total number of network errors per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]
Frame	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
in Kbytes/s	The total number of network interface bytes input per second in kilobytes [Sequential = AVG Non-Sequential = SUM]
in packets/s	The total number of network input packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]
mpid	The identifier of the memory pool that the LPAR belongs to [Non-Sequential = ID]
out Kbytes/s	The total number of network interface bytes output per second in kilobytes [Sequential = AVG Non-Sequential = SUM]
out packets/s	The total number of network output packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]
partition_group_id	The LPAR group that the LPAR is a member of [Non-Sequential = ID]

partition_name	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
shared_pool_id	The identifier of the shared pool of physical processors that the LPAR is a member of [Non-Sequential = ID]
total Kbytes/s	The total number of network interface bytes input and output per second in kilobytes [Sequential = AVG Non-Sequential = SUM]
total packets/s	The total number of network packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]

Parameter Hierarchy

Class:	Network
Subclass:	by Interface
IT Resource Name:	/TeamQuest/System/systemname/Network
TeamQuest Table Name:	Network.by Interface
Open Table Name:	NETBYINTERFACE
Resource:	interface0, interface1, ...
Statistic Name:	
collisions/s	The number of network collisions per second on Carrier Sense Multiple Access (CSMA) interfaces. This value does not include Ethernet interfaces. [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-errs.rpt
idrops/s	The number of network interface input drops per second. This statistic is not available and will be reported as zero. [Sequential = AVG Non-Sequential = SUM]
ifspeed	The network interface line speed in megabits per second. This statistic is not available and will be reported as zero. [Sequential = LST Non-Sequential = SUM] View Report: /report/aix/network/net-bits.rpt
in bytes/s	The number of network interface bytes input per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-bytes.rpt
in errors/s	The number of network input errors per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-errs.rpt
in packets/s	The number of network input packets per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-pkts.rpt

Mbits/s	The network interface bit count in megabits per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-bits.rpt
out bytes/s	The number of network interface bytes output per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-bytes.rpt
out errors/s	The number of network output errors per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-errs.rpt
out packets/s	The number of network output packets per seconds [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-pkts.rpt
Class:	Network
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	Network.Summary
Open Table Name:	NETSUM
Statistic Name:	
errors/s	The total number of network errors per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]
in Kbytes/s	The total number of network interface bytes input per second in kilobytes [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/sum-bytes.rpt
in packets/s	The total number of network input packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-sum.rpt
out Kbytes/s	The total number of network interface bytes output per second in kilobytes [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/sum-bytes.rpt
out packets/s	The total number of network output packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/net-sum.rpt

total Kbytes/s	The total number of network interface bytes input and output per second in kilobytes [Sequential = AVG Non-Sequential = SUM]
total packets/s	The total number of network packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]

Note: *NFS statistics are only collected if TeamQuest Manager is running under a privileged user ID.*

Class:	NFS
Subclass:	Client
IT Resource Name:	/TeamQuest/System/systemname/Network
TeamQuest Table Name:	NFS.Client
Open Table Name:	NFSCLI
Statistic Name:	

badcalls/s	The total number of Network File System (NFS) calls per second rejected from the client [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/nfs/client.rpt
calls/s	The total number of NFS calls sent by the client [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/nfs/client.rpt
gets/s	The total number of times per second an NFS client handle was received [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/nfs/client.rpt

Class:	NFS
Subclass:	Server
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFS.Server
Open Table Name:	NFSSERV
Statistic Name:	
badcalls/s	The total number of NFS calls per second rejected by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/nfs/server.rpt
calls/s	The total number of NFS calls per second received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/nfs/server.rpt
Class:	NFSv2
Subclass	Client
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFSv2.Client
Open Table Name:	NFSV2CLI
Statistic Name:	
calls/s*	The number of NFS version 2 calls per second sent by the client [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/nfs/clnt-v2s.rpt
Class:	NFSv2
Subclass:	Client
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFSv2.Client
Open Table Name:	NFSV2CLI
Resource:	create, getattr, link, lookup, mkdir, null, read, readdir, readlink, remove, rename, rmdir, root, setattr, statfs, symlink, write, writecache
Statistic Name:	
reqs/s	The number of NFS version 2 requests per second sent by the client [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/nfs/clnt-v2.rpt

Class: NFSv2
 Subclass: Server
 IT Resource Name: /TeamQuest/System/*systemname*/Network
 TeamQuest Table Name: NFSv2.Server
 Open Table Name: NFSV2SERV
 Statistic Name:
 calls/s* The number of NFS version 2 calls per second received by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/aix/network/nfs/srvr-v2s.rpt

Class: NFSv2
 Subclass: Server
 IT Resource Name: /TeamQuest/System/*systemname*/Network
 TeamQuest Table Name: NFSv2.Server
 Open Table Name: NFSV2SERV
 Resource: create, getattr, link, lookup, mkdir, null, read, readdir, readlink,
 remove, rename, rmdir, root, setattr, statfs, symlink, write, writecache
 Statistic Name:
 reqs/s The number of NFS version 2 requests per second received by the
 server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/aix/network/nfs/srvr-v2.rpt

Class: NFSv3
 Subclass: Client
 IT Resource Name: /TeamQuest/System/*systemname*/Network
 TeamQuest Table Name: NFSv3.Client
 Open Table Name: NFSV3CLI
 Statistic Name:
 calls/s* The number of NFS version 3 calls per second sent by the client
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/aix/network/nfs/clnt-v3s.rpt

Class: NFSv3
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Client
Open Table Name: NFSV3CLI
Resource: access, commit, create, fsinfo, fsstat, getattr, link, lookup, mkdir, mknod, null, pathconf, read, readdir, readdir+, readlink, remove, rename, rmdir, setattr, symlink, write
Statistic Name:
 reqs/s The number of NFS version 3 requests per second sent by the client
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/aix/network/nfs/clnt-v3.rpt

Class: NFSv3
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Server
Open Table Name: NFSV3SERV
Statistic Name:
 calls/s* The number of NFS version 3 calls per second received by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/aix/network/nfs/srvr-v3s.rpt

Class: NFSv3
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Server
Open Table Name: NFSV3SERV
Resource: access, commit, create, fsinfo, fsstat, getattr, link, lookup, mkdir, mknod, null, pathconf, read, readdir, readdir+, readlink, remove, rename, rmdir, setattr, symlink, write
Statistic Name:
 reqs/s The number of NFS version 3 requests per second received by the
 server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/aix/network/nfs/srvr-v3.rpt

Note: *RPC statistics are only collected if TeamQuest Manager is running under a privileged user ID.*

Class:	RPC
Subclass:	Client.Connectionless
IT Resource Name:	/TeamQuest/System/systemname/Network
TeamQuest Table Name:	RPC.Client.Connectionless
Open Table Name:	RPCCLICONNLESS
Statistic Name:	
badcalls/s	<p>The number of connectionless RPC calls per second rejected by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-cl.rpt</p>
badxid/s	<p>The number of times per second a reply from a server was received that did not correspond to any outstanding connectionless RPC call</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-cl.rpt</p>
calls/s	<p>The total number of connectionless RPC calls per second sent by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-cl.rpt</p>
newcred/s	<p>The number of times per second connectionless RPC authentication information had to be refreshed by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-cl.rpt</p>
retrans/s	<p>The number of times per second a connectionless RPC call had to be retransmitted by the client due to a timeout while waiting for a reply from the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-cl.rpt</p>
timeout/s	<p>The number of times per second a connectionless RPC call timed out while waiting for a reply from the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-cl.rpt</p>

Class:	RPC
Subclass:	Client.Connection Oriented
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	RPC.Client.Connection Oriented
Open Table Name:	RPCCLICONNORIENTED
Statistic Name:	
badcalls/s	<p>The number of connection-oriented RPC calls per second rejected from the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-co.rpt</p>
badxid/s	<p>The number of times per second a reply from a server was received that did not correspond to any outstanding connection-oriented RPC call</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-co.rpt</p>
calls/s	<p>The total number of connection-oriented RPC calls per second sent by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-co.rpt</p>
newcred/s	<p>The number of times per second connection-oriented RPC authentication information had to be refreshed by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-co.rpt</p>
timeout/s	<p>The number of times per second a connection-oriented call timed out while waiting for a reply from the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/network/rpc/clnt-co.rpt</p>

Class:	RPC
Subclass:	Server.Connectionless
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	RPC.Server.Connectionless
Open Table Name:	RPCSERVCONNLESS
Statistic Name:	
badcalls/s	<p>The number of connectionless RPC calls per second rejected by the server. This is the sum of badlens/s and xdrcalls/s. [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-cl.rpt</p>
badlens/s	<p>The number of connectionless RPC calls per second received by the server with a length shorter than a minimum-sized RPC call [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-cl.rpt</p>
calls/s	<p>The number of connectionless RPC calls per second received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-cl.rpt</p>
nullrecvs/s	<p>The number of times per second a connectionless RPC call was not available when it was thought to be received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-cl.rpt</p>
xdrcalls/s	<p>The number of connectionless RPC calls per second by the server whose header could not be External Data Representation (XDR) decoded [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-cl.rpt</p>

Class:	RPC
Subclass:	Server.Connection Oriented
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	RPC.Server.Connection Oriented
Open Table Name:	RPCSERVCONNORIENTED
Statistic Name:	
badcalls/s	<p>The number of connection-oriented RPC calls per second rejected by the server. This is the sum of badlens/s and xdrcalls/s. [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-co.rpt</p>
badlens/s	<p>The number of connection-oriented RPC calls per second received by the server with a length shorter than a minimum-sized RPC call [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-co.rpt</p>
calls/s	<p>The number of connection-oriented RPC calls per second received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-co.rpt</p>
nullrecvs/s	<p>The number of times per second a connection-oriented RPC call was not available when it was thought to be received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-co.rpt</p>
xdrcalls/s	<p>The number of connection-oriented RPC calls per second by the server whose header could not be XDR decoded [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/network/rpc/srvr-co.rpt</p>

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
net interval	The number of seconds elapsed between two data samples of the Network Agent [Sequential = SUM Non-Sequential = ID]
tqbnp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqbnp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

6.4. Workload Manager Statistics

Workload manager statistics are stored in the TeamQuest performance database tables by the Process-Workload Agent. The workload manager statistics are available when the workload manager (WLM) workload set is turned on.

Table Field Hierarchy

Class:	AIX
Subclass:	WLM
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /workload/ <i>workloadset</i> / <i>workload</i>
TeamQuest Table Name:	AIX.WLM
Open Table Name:	AIXWLM
Collection interval:	N/A
Default retention:	1 year
Table type:	Event

Statistic Name	Description
bio_hardmax	The hard maximum limit that is allocated for the WLM disk I/O resource [Non-Sequential = SUM]
bio_min	The minimum value that is allocated for the WLM disk I/O resource [Non-Sequential = SUM]
bio_shares	The number of shares that are allocated for the WLM disk I/O resource. If a given resource is not managed by the WLM, a default value of -1 is used. [Non-Sequential = ID]
bio_softmax	The soft maximum limit that is allocated for the WLM disk I/O resource [Non-Sequential = SUM]
class	The class information that is passed to the WLM [Non-Sequential = ID]
cpu_hardmax	The hard maximum limit allocated for the WLM CPU resource [Non-Sequential = SUM]
cpu_min	The minimum value that is allocated for the WLM CPU resource [Non-Sequential = SUM]
cpu_shares	The number of shares that are allocated for the WLM CPU resource. If a given resource is not managed by the WLM, a default value of -1 is used. [Non-Sequential = ID]
cpu_softmax	The soft maximum limit that is allocated for the WLM CPU resource [Non-Sequential = SUM]
mem_hardmax	The hard maximum limit allocated for the WLM memory resource [Non-Sequential = SUM]
mem_min	The minimum value that is allocated for the WLM memory resource [Non-Sequential = SUM]

mem_shares	The number of shares that are allocated for the WLM memory resource. If a given resource is not managed by the WLM, a default value of -1 is used. [Non-Sequential = ID]
mem_softmax	The soft maximum limit that is allocated for the WLM memory resource [Non-Sequential = SUM]
System	The system name as assigned by the operating system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
tier	The tier number for the class [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]

6.5. Workload Statistics

Workload statistics are maintained in the TeamQuest performance database by the Process-Workload Agent. The statistics are classified by the hierarchy of key names.

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
tqwarp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqwarp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]
warp interval	The number of seconds elapsed between two data samples of the Process-Workload Agent [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

Parameter Hierarchy

Class:	Workload
Subclass:	by Workload
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /workload/ <i>workloadset</i> / <i>workload</i>
TeamQuest Table Name:	Workload.by Workload
Open Table Name:	WLBYWORKLOAD
Workload Set:	WLS1, WLS2, ...
Workload:	WL1, WL2, ...
Statistic Name:	
%cpu	<p>The percentage of total CPU consumed by the workload. Total CPU time is the value of the sampling interval multiplied by the number of CPUs on the system. Thus, if the sum of %cpu for all workloads is less than 100%, some CPUs must have been idle for some time during the sampling interval. For shared LPARs, this value represents the percentage of the entitled capacity used. For shared LPARs that are uncapped, the sum of this value across all workloads in a workload set can be greater than 100%. [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/workload/pct-cpu.rpt</p>
avgmem	<p>The cumulative swappable process image size in kilobytes of all of the running processes in the workload at the end of the sampling interval [Sequential = AVG Non-Sequential = SUM]</p>
etime	<p>The sum of the elapsed times in seconds of all of the processes in the workload. Dividing this number by the number of processes in the workload (pongoing + pcomplete) equals the average time a process in the workload existed during the sampling interval. [Sequential = SUM Non-Sequential = SUM]</p>
lioch	<p>The number of logical characters in kilobytes transferred by the workload during the sampling interval [Sequential = SUM Non-Sequential = SUM] View Report: /report/aix/workload/lioch.rpt</p>
majflt	<p>The number of major page faults generated by the workload for processes that were active at the end of the sampling interval. A major page fault is a page fault that requires I/O. [Sequential = SUM Non-Sequential = SUM] View Report: /report/aix/workload/maj-flt.rpt</p>
pcomplete	<p>The number of processes completed in the sampling interval. For process data, the same number is called cproc. [Sequential = SUM Non-Sequential = SUM] View Report: /report/aix/workload/num-proc.rpt</p>

pongoing	<p>The number of processes running at the end of the sampling interval. In process data, the same number is called oproc. The number of processes in a workload could be derived by adding pongoing and pcomplete. This sum is called nproc in process data. [Sequential = LST Non-Sequential = SUM] View Report: /report/aix/workload/num-proc.rpt</p>
prss	<p>An estimate of the resident set size in kilobytes of private memory occupied by all of the running processes in the workload at the end of the sampling interval. The value is derived from memory integrals kept by the kernel. [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/workload/rss.rpt</p>
pstart	<p>The number of processes started in the sampling interval. In process data, this number is called sproc. [Sequential = SUM Non-Sequential = SUM] View Report: /report/aix/workload/num-proc.rpt</p>
record_count	<p>The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record. [Sequential = AVG Non-Sequential = SUM]</p>
reduction_name	<p>The name of reduction rule [Sequential = ID Non-Sequential = ID]</p>
reduction_source	<p>The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H. [Sequential = ID Non-Sequential = ID]</p>
rss	<p>The resident set size (real memory) in kilobytes of all processes running at the end of the interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>. [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/workload/rss.rpt</p>
srss	<p>An estimate of the resident set size in kilobytes of shared memory occupied by all of the running processes in the workload at the end of the sampling interval. The value is derived from memory integrals kept by the kernel. [Sequential = AVG Non-Sequential = SUM] View Report: /report/aix/workload/rss.rpt</p>

syscpu	<p>The system CPU time in seconds used by this workload. System CPU time is the time spent in kernel mode (for example, the time spent in executing system calls, paging, and so on).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/workload/sys-cpu.rpt</p>
threads	<p>The number of threads at the end of the sampling interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/workload/threads.rpt</p>
totcpu	<p>The total CPU time in seconds used by the workload. This value is the same as the sum of usrcpu and syscpu.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/workload/cpu-util.rpt</p>
usrcpu	<p>The user CPU time in seconds used by the workload. User CPU time is the time the CPU spent running in user mode.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/aix/workload/user-cpu.rpt</p>

6.6. LPAR Configuration Statistics

The following statistics are stored in the TeamQuest performance database tables by the System Activity Agent.

Table Field Hierarchy

Class:	IBM POWER
Subclass:	LPAR
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	IBM POWER.LPAR
Open Table Name:	IBMPOWERLPAR
Collection interval:	N/A
Default retention:	1 year
Table type:	Event

Statistic Name	Description
active_cpus_in_pool	The current number of active physical CPUs in the shared processor pool being used by the LPAR [Non-Sequential = SUM]
active_physical_cpus_in_system	The current number of active physical CPUs in the system containing the LPAR [Non-Sequential = SUM]
AME	The Active Memory Expansion status of the partition [Non-Sequential = ID]
AMS	The Active Memory Sharing status of the partition [Non-Sequential = ID]
capacity_increment	The granularity at which changes to the entitled capacity can be made. A value in whole multiples indicates a dedicated LPAR. [Non-Sequential = NON]
dedicated_donating	Indicates if the partition is able to donate dedicated processor capacity [Non-Sequential = ID]
entitled_capacity	The number of processing units the LPAR is entitled to receive [Non-Sequential = SUM]
installed_physical_cpus_in_system	The current number of physical CPUs installed on the system containing the LPAR [Non-Sequential = SUM]
maximum_capacity	The maximum number of processing units the LPAR was defined to ever have. Entitled capacity can be increased up to this value. [Non-Sequential = SUM]
maximum_memory	The maximum possible amount of memory in megabytes [Non-Sequential = SUM]
maximum_physical_cpus_in_system	The maximum possible number of physical CPUs in the system containing the LPAR [Non-Sequential = MAX]

maximum_virtual_cpus	The maximum possible number of CPUs (virtual engines) [Non-Sequential = SUM]
minimum_capacity	The minimum number of processing units the LPAR was defined to ever have. Entitled capacity can be reduced down to this value. [Non-Sequential = SUM]
minimum_memory	The minimum memory in megabytes the LPAR was defined to ever have [Non-Sequential = SUM]
minimum_virtual_cpus	The minimum possible number of CPUs the LPAR was defined to ever have [Non-Sequential = SUM]
mode	Indicates whether the LPAR processor capacity is capped, or uncapped and allowed to consume idle cycles from the shared pool. A dedicated LPAR is implicitly capped. [Non-Sequential = ID]
mpid	The identifier of the memory pool that the LPAR belongs to [Non-Sequential = ID]
online_memory	The amount of memory in megabytes currently online [Non-Sequential = SUM]
online_virtual_cpus	The number of CPUs (virtual engines) currently online [Non-Sequential = SUM]
partition_group_id	The LPAR group that the LPAR is a member of [Non-Sequential = ID]
partition_name	The logical partition name as assigned by the management entity [Non-Sequential = ID]
partition_number	The number of the logical partitions assigned by the management entity [Non-Sequential = ID]
physical_cpu_percentage	The fractional representation relative to whole physical CPUs that these LPAR virtual CPUs equate to. This is a function of Entitled Capacity/Online CPUs. Dedicated LPARs would have a Physical CPU Percentage of 100%. A 4-way virtual with Entitled Capacity of 2 processor units would have a Physical CPU Percentage of 50%. [Non-Sequential = SUM]
shared_pool_id	The identifier of the shared pool of physical processors that the LPAR is a member of [Non-Sequential = ID]
smt	Indicates whether simultaneous multi-threading (SMT) is enabled. This value is SMT- <i>n</i> if SMT is enabled, where <i>n</i> is the number of threads per core. [Non-Sequential = ID]
System	The system name as assigned by the operating system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]

system_name	The name of the physical system as displayed by the management entity [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
type	Indicates whether the LPAR is using dedicated or shared CPU resources [Non-Sequential = ID]
unallocated_capacity	The number of processor units currently unallocated in the shared processor pool being used by the LPAR [Non-Sequential = NON]
Unlicensed Physical CPUs in System*	The current number of physical CPUs installed in the system containing the LPAR which are inactive because they are unlicensed. The CPUs may be available for activation by Capacity on Demand (COD) functions. [Non-Sequential = SUM]
variable_capacity_weight	The priority weight assigned to the LPAR, which controls how extra (idle) capacity is allocated to it. A weight of -1.0 indicates a soft maximum cap is in place. [Non-Sequential = NON]

Table Field Hierarchy

Class:	IBM POWER
Subclass:	LPAR FRAME
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	IBM POWER.LPAR FRAME
Open Table Name:	IBMPowerLPARFRAME
Collection interval:	Based on the primary aggregation set
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 8-hour intervals
Table type:	Performance

Note: *The IBM POWER.LPAR FRAME default table retentions should closely match the database aggregation set retentions to avoid generating charts with unexpected granularities. The IBM POWER.LPAR FRAME table is used in derived table definitions to join the Frame field with other IBM POWER tables that have retentions based on the database aggregation set retentions.*

Statistic Name	Description
Interval	The desired collection interval [Sequential = SUM Non-Sequential = AVG]
Frame	The physical system name as displayed by the management entity [Sequential = LST Non-Sequential = ID]
mpid	The identifier of the memory pool that the LPAR belongs to [Non-Sequential = ID]
partition_group_id	The LPAR group that the LPAR is a member of [Non-Sequential = ID]
partition_name	The logical partition name as assigned by the management entity [Sequential = LST Non-Sequential = ID]
shared_pool_id	The identifier of the shared pool of physical processors that the LPAR is a member of [Non-Sequential = ID]
System	The system name as assigned by the operating system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Table Field Hierarchy

Class:	AIX
Subclass:	WPAR
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	AIX.WPAR
Open Table Name:	AIXWPAR
Collection interval:	N/A
Default retention:	1 year
Table type:	Event

Statistic Name	Description
cpu_limit	The percentage of the partition's CPU cycles the WPAR is entitled to receive [Non-Sequential = ID]
entitled_capacity	The number of processors the partition hosting the WPAR is entitled to receive [Non-Sequential = SUM]
mem_limit	The percentage of the partition's memory the WPAR is entitled to receive [Non-Sequential = ID]
System	The name of the global environment which contains the WPAR [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
wpar_name	The name of the WPAR in the global environment [Non-Sequential = ID]

6.7. Process Statistics

The Process-Workload Agent collects process data from the operating system and processes accounting files. The Process-Workload Agent calculates the usage of every process in a given interval, applies reduction definitions to each process, and stores the reduced process data. It also applies workload definitions to the reduced process data and stores system resource usage by workload.

Reduced Process Records

The Process-Workload Agent collects process data and reduces the data according to the user-defined reduction definitions. A reduction definition may cause multiple processes to be merged into a single process record. Thus, a process record contains data about one or more processes. When you are looking at the resource usage numbers, it is important to know how many processes a process record actually represents. The `nproc` data item indicates exactly how many processes each process record is representing. When a process record is representing more than one process, the resource usage fields such as `totcpu`, `rss`, and `pio_t` are the sum of the resource usage of the individual processes. When all of the processes do not have the same value for a field, the identifier fields such as `command`, `login`, and `gid` are set to `<Multi>`. When data for some fields is not available, the fields are set to `<N/A>`.

Disabling Reduction Definitions

If you want to look at the details of every individual process and do not wish to have merged process records, you must disable reduction processing by making all reduction sets inactive. However, with reduction processing disabled, more records have to be stored and more disk space is needed. For information on disabling reductions, see the *TeamQuest Performance Software Administration Guide*.

Process Data with Process Accounting Turned Off

If process accounting is turned off, the process data is incomplete as data about completed processes is not available. In this case, the process data only shows a portion of the activity in the sampling interval. To find out whether process accounting is turned on or off, look at the `cproc` field of all of the process records. If the `cproc` field indicates all zeros, it means that processes were not completed in that sampling interval and that the process accounting is turned off.

When process accounting is not turned on, the process record, `<Other>` includes the CPU time for processes that completed during the interval.

For more information on process accounting, see the *TeamQuest Performance Software Administration Guide*.

Retrieving Hardware Configuration Information

The Process-Workload Agent retrieves hardware configuration information. The information is stored upon startup and once-a-day in the HINV.Summary, HINV.Devices, and HINV.FileSystem table files of the TeamQuest performance database. The information is also stored if the agent detects a change in configuration.

Notes:

- *The statistics reported in the AIX.Process table on a shared logical partition are relative to the entitlement of the partition.*
- *The statistics reported in the AIX.Process table on a dedicated logical partition are relative to the configured capacity of the partition.*
- *For more information on partition entitlement, shared partitions, or dedicated partitions, see the IBM PowerVM Virtualization documentation.*

Table Field Hierarchy

Class:	AIX
Subclass:	Process
IT Resource Name:	/TeamQuest/System/systemname/Process
TeamQuest Table Name:	AIX.Process
Open Table Name:	AIXPROC
Collection interval:	Based on the primary aggregation set
Default retention:	1 day
Table type:	Performance

Note: *The collection interval is also dependent on the Processes Only setting in the configuration file for the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.*

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgmem_t	The size of the swappable process image in kilobytes. If a process ends within the interval, the number is unavailable and is marked as <N/A>. [Sequential = AVG Non-Sequential = SUM]
btime	The start time of the process. For process records representing more than one process, this field shows the earliest of the start times. [Sequential = FST Non-Sequential = ID]

Class	<p>The AIX Workload Manager (WLM) class to which the process belongs. This field will be marked <N/A> if the WLM is not running or the class cannot be obtained.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
command	<p>The command name of the process. If a process starts and ends within the same interval, only up to 8 characters of the command name can be displayed. Otherwise, up to 14 characters are displayed. Therefore, an “automountd” process may appear as “automoun” if it starts and ends within the same interval. In such cases, you may want to set up your workload, reduction, or filter definitions to catch both the command names as in the following:</p> <p>command = {“automoun”, “automountd”}.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
cproc	<p>The number of processes completed in the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
etime	<p>The elapsed time in seconds for the current interval. This number tells how long a process existed in the current interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
etime_t	<p>The total elapsed time in seconds. This number tells how long a process existed since it started. For a single process, this is the same as the value of the Time field minus the value of the btime field.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
fullcmd	<p>The full command string, including arguments, for the process. If a process begins and ends in the sampling interval, the field is unavailable and is marked as <N/A>. This is an important consideration when using a workload, reduction, or filter definition. The number of characters that are stored is determined by the Command Length setting of the Process-Workload Agent configuration file in TeamQuest Manager. You can also have either the first or the last <i>N</i> characters of the command displayed. The setting is specified by the Command Orientation setting of the Process-Workload Agent configuration file in TeamQuest Manager. The limit of characters from the operating system data source is 16,383. For a description of how this may affect collection when the LAST Command Orientation is configured for the Process-Workload Agent, see the <i>TeamQuest Performance Software Administration Guide</i>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
gid	<p>The real group identifier of the owner of the process</p> <p>[Sequential = ID Non-Sequential = ID]</p>
group	<p>The group name of the owner of the process. This field is derived from gid.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Interval	<p>The expected data sampling interval</p> <p>[Sequential = SUM Non-Sequential = ID]</p>

lioch	The number of logical characters in kilobytes transferred in the current interval. The number reported represents only completed processes during the sampling interval. [Sequential = SUM Non-Sequential = SUM]
lioch_t	The total number of logical characters in kilobytes transferred since the process started. The number reported represents only the completed processes during the sampling interval. [Sequential = LST Non-Sequential = SUM]
login	The login name of the owner of the process. This field is derived from uid. [Sequential = ID Non-Sequential = ID]
majflt	The number of major page faults generated in the current interval. A major page fault is a page fault that requires I/O. If a process ends in the sampling interval, the number is unavailable and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
majflt_t	The total number of major page faults generated since the process started. A major page fault is a page fault that requires I/O. If the process starts and ends within the same interval, the number is unavailable and is marked as <N/A>. [Sequential = LST Non-Sequential = SUM]
maxrss	The maximum resident set size in kilobytes for the process. If a process ends within the sampling interval, the value is unavailable and is marked as <N/A>. [Sequential = MAX Non-Sequential = SUM]
nproc	The number of processes that the process record represents. In a reduced record, it is the number of processes that were merged together to form a single process record. When no time consolidation is applied to the process data, this number should equal the sum of cproc and oproc. When consolidating over time, the number represents the average number of processes in the process record for the consolidated period. An average is used because processes can move in and out and between reduced records from interval to interval. [Sequential = AVG Non-Sequential = SUM]
oproc	The number of ongoing processes at the end of the interval [Sequential = LST Non-Sequential = SUM]
pctcpu	The percentage of total available CPU time the process used in the current sampling interval. For shared LPARs, this value represents the percentage of the entitled capacity used. For shared LPARs that are uncapped, the sum of this value across all processes can be greater than 100%. [Sequential = AVG Non-Sequential = SUM]
pid	The process identifier number. If a process starts and ends within an interval, the number is unavailable and is marked as <N/A>. [Sequential = ID Non-Sequential = ID]

ppid	<p>The numerical identifier of the parent of a process. If a process starts and ends within an interval, the number is unavailable and is marked as <N/A>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
pri	<p>The priority of the process. Higher numbers mean lower priority. If a process ends within an interval, the number is unavailable and is marked as <N/A>. If the process record represents more than one process, the priority is an average of the constituent running processes' priority values.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
prss	<p>An estimate of the resident set size in kilobytes of private memory for the process at the end of the sampling interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
redname	<p>The reduction name of the process record. If a process did not match any of the reduction definitions, then it would not be reduced and will not have a reduction name.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
rss	<p>The resident set size (real memory size) of the process at the end of the interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
sproc	<p>The number of processes started in the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
srss	<p>The resident set size in kilobytes of shared memory for the process at the end of the sampling interval. The value is derived from memory integrals kept by the kernel.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
syscpu	<p>The system CPU time in seconds for the current interval. System CPU time is the time the CPU spent running in kernel mode (for example, the time spent in executing system calls, paging, and so on). If an application is taking a lot of syscpu time, you may want to optimize the use of system calls (for example, use a larger block size for I/O).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
syscpu_t	<p>The total system CPU time in seconds</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
threads	<p>The number of threads associated with the process at the end of the interval. A thread is a dynamic object that represents a control point in a process and executes a sequence of instructions.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totcpu	The total CPU time in seconds used in the current interval. This number is the same as the sum of usrcpu and syscpu. [Sequential = SUM Non-Sequential = SUM]
totcpu_t	The total CPU time (user + system) in seconds used by the process since it started. This number is the same as the sum of usrcpu_t and syscpu_t. [Sequential = LST Non-Sequential = SUM]
tty	The controlling terminal identifier in dev_t format. For the processes without a controlling terminal, this field will contain a -1. [Sequential = ID Non-Sequential = ID]
ttyname	The controlling terminal for the process. It is a device name without the /dev/ prefix. This is derived from tty. For the processes without a controlling terminal, this field will contain a question mark (?). [Sequential = ID Non-Sequential = ID]
uid	The real user id of the process owner [Sequential = ID Non-Sequential = ID]
usrcpu	The user CPU time in seconds for the current interval. User CPU time is the time the CPU spent running user mode. If an application is taking a lot of usrcpu time, you should try to optimize the code, if possible. [Sequential = SUM Non-Sequential = SUM]
usrcpu_t	The total user CPU time in seconds since the start of the process [Sequential = LST Non-Sequential = SUM]
WPAR	The name of the WPAR in the global environment [Sequential = ID Non-Sequential = ID]
Workload	<p>The workload set and the workload associated with the process. When the Process-Workload Agent stores the process record, this field is blank. When process records are reported, the workload can be evaluated and is shown in the report.</p> <p>This field is available for reporting only when using TeamQuest Analyzer and TeamQuest tView.</p> <p>Workload evaluation takes place when data is retrieved from the database, based on workload sets defined in the database where the data is stored. Workload sets reported in the Workload field do not have to be active.</p> <p>For more information on workload evaluation, see the <i>TeamQuest Analyzer User Guide</i> or the <i>TeamQuest Performance Software Command Line Interfaces Reference Manual</i>. [Sequential = ID Non-Sequential = ID]</p>

workload:*wlsname* There is one field for each *wlsname* (Workload Set Name). The value for this field shows the name of the workload to which the process belongs. If a process belongs to none of the workloads, it will display the workload name “OTHER.”

This field is available for reporting only when using TeamQuest View or TeamQuest cView.
[Sequential = ID Non-Sequential = ID]

6.8. Hardware Inventory Statistics

The hardware inventory statistics listed in this section are stored in the TeamQuest performance database tables by the Process-Workload Agent.

The Process-Workload Agent retrieves hardware configuration information. The information is stored upon startup and once-a-day in the HINV.Summary, HINV.CPUModel, HINV.CPU Thread Speeds, HINV.Devices, and HINV.FileSystem table files of the TeamQuest performance database. The information is also stored if the agent detects a change in configuration.

Note: *The storage of hardware inventory records depends on the Hardware Inventory setting in the configuration file of the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.*

Table Field Hierarchy

Class:	HINV
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.Summary
Open Table Name:	HINVSUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
core_multi_thread	The status or ability of the processor to support multiple independent threads. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cores_per_chip	The number of cores on an individual chip. This statistic is not available for IBM AIX POWER7 and the value is displayed as <N/A>. [Non-Sequential = ID]
cpu_chips	The number of CPU chips. This statistic is not available for IBM AIX POWER7 and the value is displayed as <N/A>. [Non-Sequential = ID]

cpu_count	The current number of active physical CPUs in the system. For LPAR capable systems, the current number of active physical CPUs in the system containing the LPAR. [Non-Sequential = ID]
cpu_speed	The speed of the processor in MHz or GHz [Non-Sequential = ID]
cpu_type	The basic instruction set architecture of the current system [Non-Sequential = ID]
logical_cpu_count	The number of logical processors. This value represents the number of processors as seen by the operating system. This value is influenced by the virtual processor and simultaneous multithreading (SMT) settings. [Non-Sequential = ID]
mem_size	The size of configured random access memory in kilobytes, where 1 kilobyte = 1,024 bytes [Non-Sequential = ID]
memory	The size of configured random access memory in megabytes, where 1 megabyte = 1,048,576 bytes [Non-Sequential = ID]
memory_size	The size of configured random access memory in megabytes or gigabytes [Non-Sequential = ID]
model	Name of the hardware implementation or platform [Non-Sequential = ID]
os_release	The name and level of this implementation of the operating system [Non-Sequential = ID]
pagesize	The size of a page of memory [Non-Sequential = ID]
partition_type	The partition type of the system. The value indicates the system hypervisor type, guest type, logical partition type, zone type, or logical domain type. If the system does not have a partition type, this field will be blank. [Non-Sequential = ID]
serial	The hardware specific serial number of the physical machine [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_identifier	Information used to identify the system [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]

timezone	The time zone where the data was collected [Non-Sequential = ID]
TQLevel	The level of TeamQuest Manager [Non-Sequential = ID]

The HINV.CPUModel table stores best-match, relative performance data about the system configuration. This table is created by the Hardware Inventory Agent (**tqhinv**) to map physical hardware to a CPU model that describes performance in relative terms. This table is not created for any virtualized system. It is populated for physical systems only. It is not populated for VMware guests, Hyper-V guests, Solaris LDOMs, Solaris guest LDOMs, KVM guests, and Linux on POWER systems.

Table Field Hierarchy

Class:	HINV
Subclass:	CPUModel
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	HINV.CPUModel
Open Table Name:	HINVCPUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
cpu_chips	The number of CPU chips or sockets [Non-Sequential = ID]
cpu_confidence	The percentage of confidence in the correctness of the CPU match based on model, frequency, and configuration (chips, cores, threads) [Non-Sequential = SUM]
cpu_cores	The number of CPU cores or processors on an individual CPU chip [Non-Sequential = ID]
cpu_name	The name of the selected CPU [Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Non-Sequential = ID]
cpu_speed	The speed of the processor in megahertz (MHz) or gigahertz (GHz) [Non-Sequential = ID]
cpu_threads	The number of CPU threads on an individual CPU core or processor [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]

system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
user_override	The user override status of the default TeamQuest generated CPU match. This field is not currently used and should appear as 0. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	CPU Thread Speeds
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	HINV.CPU Thread Speeds
Open Table Name:	HINVCPUTHREADSPEEDS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
speed_up_factor	The performance improvement when there are multiple active threads per core, compared to when there is only one active thread per core [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
thread_number	The number of active threads [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	Devices
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.Devices
Open Table Name:	HINVDEVS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
class	The device classification: controller, disk, or tape [Non-Sequential = ID]
controller	The device path indicator which defines a connection to another device [Non-Sequential = ID]
lun_id	The globally unique Logical Unit Number (LUN) identifier for Storage Area Network (SAN) based disk devices. This field is blank for non-SAN based disk devices, CD-ROM drives, tape drives, and so on. [Non-Sequential = ID]
name	The unique identifier for this device [Non-Sequential = ID]
name2	The alternate device name. This field may be blank. [Non-Sequential = ID]
product	The product identifier. This field may be blank. [Non-Sequential = ID]
revision	The revision level for the product. This field may be blank. [Non-Sequential = ID]
rpm	The speed at which the media spins. If an actual value cannot be obtained for the device, a default value of 7,200 is used. [Non-Sequential = ID]
sequence	The sequence number of the device [Non-Sequential = ID]
swap	A true or false statement which indicates whether or not a swap file exists on the device. [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]

Time	The timestamp of the data sample [Non-Sequential = ID]
vendor	The name of the device vendor. This field may be blank. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	FileSystem
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.FileSystem
Open Table Name:	HINVFILESYS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
BlkSize	The size of a block on the file system [Non-Sequential = ID]
Device	The path for the device on which the file system is mounted [Non-Sequential = ID]
Name	The unique identifier for the file system [Non-Sequential = ID]
Source	The source physical disk or logical volume of the file system. This field is always blank for this platform. [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
TotBlks	The total number of blocks on the file system [Non-Sequential = ID]
TotFiles	The maximum total number of files, as represented by inodes, possible on the file system. Some inodes may be used for entities other than visible files. [Non-Sequential = ID]
TotSize	The total amount of space on the file system in megabytes [Non-Sequential = ID]
Type	The type of the file system [Non-Sequential = ID]

6.9. System Log Statistics

The System Log Agent is used to collect system log messages generated by the system log daemon (syslogd). The System Log Agent stores these messages in the TeamQuest performance database for analysis and alarm reporting. The log messages are separated into four fields; the time that the message was posted, the host system from where the message was initiated, the program or user that posted the message, and the text of the message.

Table Field Hierarchy

Class:	System
Subclass:	System Log
IT Resource Name:	/TeamQuest/System/systemname/System Log
TeamQuest Table Name:	System.System Log
Open Table Name:	SYSSYSTEMLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Event_Time	The time that the message was logged to the system log [Non-Sequential = ID]
Loghost	The name of the system that logged the message [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Reporter	The name of the user or process that logged the message [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system where the log message originated. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

6.10. General Log Statistics

The General Log Agent is used to collect log messages generated by application programs. The General Log Agent stores these messages in the TeamQuest performance database for analysis. Examples include backup, security, database, and Web server applications.

Table Field Hierarchy

Class:	System
Subclass:	General Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /General Log
TeamQuest Table Name:	System.General Log
Open Table Name:	SYSGENERALLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The message type [Non-Sequential = ID]

6.11. TeamQuest Log Statistics

The following statistics are stored in the performance database tables by the TeamQuest Log Agent. The collection interval and retention periods can be modified. For more information on modifying the collection interval and retention periods, see the *TeamQuest Performance Software Administration Guide*.

Table Field Hierarchy

Class:	Service
Subclass:	TeamQuest Log
IT Resource Name:	/TeamQuest/System/systemname/TeamQuest Log
TeamQuest Table Name:	Service.TeamQuest Log
Open Table Name:	SVCTQLOG
Collection interval:	N/A
Default retention:	1 day
Table type:	Event

Statistic Name	Description
Filename	The name of the TeamQuest log file that was the source of the message text [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The log message type. This is always set to tqlog . [Non-Sequential = ID]

6.12. Derived Statistics

Some products within TeamQuest Performance Software use derived statistics to display common statistics across different platforms. The derived statistics are inserted into the performance database when databases are created. In this section, a derived statistic is marked with an asterisk (*).

You can find information on the following derived statistics:

- Workload Performance Derived Statistics
- TeamQuest On the Web Derived Statistics
- TeamQuest Alert Derived Statistics

6.12.1. Workload Performance Derived Statistics

TeamQuest Manager maintains derived statistics that use data from the System Activity Agent and the Process-Workload Agent. The workload performance reports reference these statistics. For information on workload performance reports, see the *TeamQuest View Reports Reference Manual*.

Parameter Hierarchy

Class:	Derived
Subclass:	Workload Performance.by Workload
Workload Set:	WLS1, WLS2, ...
Workload:	ALL
Statistic Name:	
%cpu*	The total percentage of CPU utilization. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/workload.rpt
Kbytes resident memory/process*	The average amount of resident memory used per process. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/workload.rpt
Kbytes virtual memory/process*	The average amount of virtual memory used per process. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/workload.rpt
lioch/s*	The number of logical characters transferred in kilobytes per second. Collected by the Process-Workload Agent.
Population (etime/interval)*	The average number of concurrent processes. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/workload.rpt

Response (etime/process)*	The elapsed time per process. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/workload.rpt
Throughput (processes/sec)*	The number of rocesses completed per second. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/workload.rpt
Total Kbytes resident memory*	The average amount of resident memory used by the workload. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/workload.rpt
Total Kbytes virtual memory*	The average amount of virtual memory used by the workload. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/workload.rpt
Class:	Derived
Subclass:	Workload Performance.Summary
Workload Set:	WLS1, WLS2, ...
Statistic Name:	
%cpu*	The total percentage of CPU utilization. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/overall.rpt
block IO r+w/s*	Disk and tape I/Os per second. Collected by the System Activity Agent. View Report: /report/aix/wkldperf/overall.rpt
Kbytes resident memory/process*	Average resident memory used per process. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/overall.rpt
Kbytes virtual memory/process*	Average virtual memory used per process. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/overall.rpt
Population (etime/interval)*	Average number of concurrent processes. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/overall.rpt
Response (etime/process)*	Elapsed time per process. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/overall.rpt
Throughput (processes/sec)*	Processes completed per second. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/overall.rpt

Total Kbytes resident memory*	Average resident memory used. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/overall.rpt
Total Kbytes virtual memory*	Average virtual memory used. Collected by the Process-Workload Agent. View Report: /report/aix/wkldperf/overall.rpt

6.12.2. TeamQuest On the Web Derived Statistics

The derived statistics used by TeamQuest On the Web include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQWeb.Summary
Statistic Name:	
avg_disk_queue_length*	The average number of requests outstanding
avg_service_time*	The average amount of time to service each transfer request for all devices in milliseconds
buffer_pct_read_cache*	The percentage of logical reads satisfied from the buffer cache
buffer_pct_write_cache*	The percentage of logical writes satisfied from the buffer cache
disk_xfers_per_sec*	The total number of read and write transfers per second for all devices
free_disk_space*	The amount of space available (not in use) on all file systems in megabytes. This measurement is taken at the end of the sampling interval and includes the space held back from normal users.
free_real_mem*	The amount of free memory available in megabytes. This measurement is taken at the end of the sampling interval.
free_swap_space*	The number of megabytes free for process swapping. This measurement is taken at the end of the sampling interval.
nfs_calls_per_sec*	The total number of NFS calls sent by the client
page_ins_per_sec*	The number of page-in requests per second
page_outs_per_sec*	The number of page-out requests per second
page_scans_per_sec*	The rate per second at which the page daemon scans pages to see if they can be freed
pct_cpu_busy*	The percentage of total CPU time the CPU was busy (not idle). This value includes the time running system code and the time running normal priority user processes.
pct_disk_busy*	The percentage of time a disk was busy servicing a transfer request
pct_sys_cpu*	The percentage of total CPU time spent in system mode
pct_usr_cpu*	The percentage of total CPU time spent running in user mode

pkt_errors_per_sec*	The total number (in + out) of network errors per second for all network interfaces
pkts_in_per_sec*	The total number of network input packets per second for all network interfaces
pkts_out_per_sec*	The total number of network output packets per second for all network interfaces
pkts_per_sec*	The total number (in + out) of network packets per second for all network interfaces
total_disk_space*	The total (used + available) amount of space on all file systems in megabytes. This measurement is taken at the end of the sampling interval and includes the space held back from normal users.
total_processes*	The number of entries currently being used in the process table. This measurement is taken at the end of the sampling interval.
total_real_mem*	The total amount of real (physical) memory in megabytes. This measurement is taken at the end of the sampling interval.
total_swap_space*	The total number of megabytes available for swapping

6.12.3. TeamQuest Alert Derived Statistics

The derived statistics used by TeamQuest Alert include the following:

Parameter Hierarchy

Class: Derived

Subclass: TQAlert.Summary

Statistic Name:

free_real_mem*	The average amount of memory available to user processes in megabytes
net_errors*	The number of network errors for all network interfaces
page_scans*	The number of pages per second scanned by the page-stealing daemon
pct_cpu_busy*	The percentage of time the CPU was not idle
pct_wio*	The percentage of total CPU time spent idle while some process is waiting for I/O completion
run_queue*	The average length of the run queue, a queue of processes in memory and runnable while the run queue is occupied
total_processes*	The total number of processes active on the system

Section 7

KVM Systems

The Libvirt Agent collects configuration and performance data for KVM hosts and virtual machines.

This section contains a listing of the statistics collected by the agent:

- Libvirt.Host Config (see 7.1)
- Libvirt.Host CPU Summary (see 7.2)
- Libvirt.Host Device (see 7.3)
- Libvirt.Host Memory Summary (see 7.4)
- Libvirt.Host Network Config (see 7.5)
- Libvirt.Host Network Filter (see 7.6)
- Libvirt.Host Numa Memory (see 7.7)
- Libvirt.Host Storage Pool Config (see 7.8)
- Libvirt.Host Volume Config (see 7.9)
- Libvirt.Virtual CPU (see 7.10)
- Libvirt.Virtual Disk (see 7.11)
- Libvirt.Virtual Disk Config (see 7.12)
- Libvirt.Virtual Machine (see 7.13)
- Libvirt.Virtual Machine Snapshot (see 7.14)
- Libvirt.Virtual Memory (see 7.15)
- Libvirt.Virtual Network Config (see 7.16)
- Libvirt.Virtual Network Interface (see 7.17)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

7.1. Libvirt.Host Config

The Libvirt.Host Config table stores configuration information for host hardware architecture.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host Config
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host Config
Open Table Name:	LIBVIRTHOSTCONFIG
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
ActiveCPUs	The number of host CPUs that are active [Non-Sequential = LST]
Cores	The number of CPU cores per socket on the host [Non-Sequential = LST]
CPUFrequency	The expected host CPU frequency in megahertz [Non-Sequential = LST]
CPUModel	The host CPU model [Non-Sequential = NON]
MemorySize	The total installed memory on the host in kilobytes [Non-Sequential = LST]
NUMACells	The total number of Non-Uniform Memory Access (NUMA) cells on the host. This field will contain a 1 for unusual NUMA topologies or uniform memory access. [Non-Sequential = LST]
Sockets	The number of CPU sockets per NUMA cell [Non-Sequential = LST]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Threads	The number of threads per CPU core on the host [Non-Sequential = LST]
Time	The timestamp of the data sample [Non-Sequential = ID]

7.2. Libvirt.Host CPU Summary

The Libvirt.Host Summary table stores CPU information for host systems.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host CPU Summary
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host CPU Summary
Open Table Name:	LIBVIRTHOSTCPUSUMY
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
%busy	The percentage of time the CPU was not idle [Sequential = AVG Non-Sequential = AVG]
%idle	The percentage of CPU time spent idle while no processes are waiting for an I/O operation to complete [Sequential = AVG Non-Sequential = AVG]
%sys	The percentage of total CPU time spent in system mode [Sequential = AVG Non-Sequential = AVG]
%usr	The percentage of total CPU time spent in user mode [Sequential = AVG Non-Sequential = AVG]
%wio	The percentage of total CPU time spent idle while some processes are waiting for an I/O operation to complete [Sequential = AVG Non-Sequential = AVG]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

7.3. Libvirt.Host Device

The Libvirt.Host Devices table stores configuration information for host hardware devices.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host Device
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host Device
Open Table Name:	LIBVIRTHOSTDEVICE
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
Capability	The description of the host device type or capabilities [Non-Sequential = ID]
Name	The name of the host device. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
Parent	The name of the parent host device for the current host device [Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

7.4. Libvirt.Host Memory Summary

The Libvirt.Host Memory Summary table stores memory information for host systems.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host Memory Summary
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host Memory Summary
Open Table Name:	LIBVIRTHOSTMEMSUMY
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name

Description

Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
buffermem	The amount of memory in megabytes in use by buffers [Sequential = AVG Non-Sequential = AVG]
cachedmem	The amount of cached memory in use in megabytes [Sequential = AVG Non-Sequential = AVG]
freemem	The amount of memory in megaybes available for user processes [Sequential = AVG Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

totalmem	The total amount of memory in megabytes [Sequential = AVG Non-Sequential = AVG]
usedmem	The amount of memory in use in megabytes [Sequential = AVG Non-Sequential = AVG]

7.5. Libvirt.Host Network Config

The Libvirt.Host Network Config table stores configuration information for host network interface devices.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host Network Config
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host Network Config
Open Table Name:	LIBVIRTHOSTNETCONFIG
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
MacAddr	The Media Access Control (MAC) address of the host network interface [Non-Sequential = ID]
Name	The name of the host network device. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
State	The state of the host network interface [Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

7.6. Libvirt.Host Network Filter

The Libvirt.Host Network Filter table stores configuration information for host firewall rules.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host Network Filter
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host Network Filter
Open Table Name:	LIBVIRTHOSTNETFILTER
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
Name	The name of the network filter. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
UUID	The universally unique identifier for the network filter [Non-Sequential = ID]

7.7. Libvirt.Host Numa Memory

The Libvirt.Host Numa Memory table stores host free memory information for each Numa cell on the host.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host Numa Memory
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host Numa Memory
Open Table Name:	LIBVIRTHOSTNUMAMEM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
FreeMemory	The free memory of the NUMA cell in gigabytes [Sequential = LST Non-Sequential = LST]
Index	The numeric index of the Non-Uniform Memory Access (NUMA) cell [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

7.8. Libvirt.Host Storage Pool Config

The Libvirt.Host Storage Pool Config table stores configuration information for host storage pools.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host Storage Pool Config
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host Storage Pool Config
Open Table Name:	LIBVIRTHOSTSPCONFIG
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
Allocation	The used portion of the storage pool in gigabytes [Non-Sequential = AVG]
Available	The available portion of the storage pool in gigabytes [Non-Sequential = AVG]
Capacity	The logical size of the storage pool in gigabytes [Non-Sequential = MAX]
Name	The name of the host storage pool device. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
Persistent	The persistence state of the host storage pool. If the value is Persistent, the storage pool will continue to exist after the host system is restarted. If the value is Transient, the storage pool will not exist after the host system is restarted. [Non-Sequential = ID]
State	The current state of the host storage pool. This value can be any of the following: Initializing pool not available Not Running Running, but not accessible Running Degraded Running Normally [Non-Sequential = ID]

System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

7.9. Libvirt.Host Volume Config

The Libvirt.Host Volume Config table stores configuration information for host storage volumes.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Host Volume Config
IT Resource Name:	/TeamQuest/System/Open Virtualization/Host/ <i>systemname</i>
TeamQuest Table Name:	Libvirt.Host Volume Config
Open Table Name:	LIBVIRTHOSTVLCONFIG
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
Allocation	The allocated size of the storage volume in gigabytes [Non-Sequential = AVG]
Capacity	The logical size of the storage volume in gigabytes [Non-Sequential = MAX]
Name	The name of the storage volume. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
StoragePool	The name of the storage pool associated with the storage volume [Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

Type	The storage volume type. This value can be any of the following: Block based volume Regular file based volume Unknown storage volume type [Non-Sequential = NON]
Volume	The unique identifier for the storage volume [Non-Sequential = ID]

7.10. Libvirt.Virtual CPU

The Libvirt.Virtual CPU table stores virtual CPU performance information for each virtual machine.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Virtual CPU
IT Resource Name:	/TeamQuest/System/Open Virtualization/Virtual Machine/ <i>virtualmachinename</i>
TeamQuest Table Name:	Libvirt.Virtual CPU
Open Table Name:	LIBVIRTVIRTUALCPU
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
%busyphys	The percentage of CPU time used by the virtual CPU during the collection interval relative to the total physical CPU resources of the host [Sequential = AVG Non-Sequential = SUM]
%busyvcpu	The percentage of CPU time used by the virtual CPU during the collection interval relative to the total virtual CPU resources of the virtual machine [Sequential = AVG Non-Sequential = AVG]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]

CPUTime	The number of CPU seconds used by the virtual CPU during the collection interval [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
PhysicalCPU	The physical CPU associated with the virtual CPU. This value is captured at the end of the sampling interval. [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
State	The current state of the virtual CPU. This value can be any of the following: Blocked Offline Running Unknown [Sequential = ID Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCPUID	The unique identifier of the virtual CPU [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine (guest) to which the data applies. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]

7.11. Libvirt.Virtual Disk

The Libvirt.Virtual Disk table stores virtual disk performance information for each virtual machine.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Virtual Disk
IT Resource Name:	/TeamQuest/System/Open Virtualization/Virtual Machine/ <i>virtualmachinename</i>
TeamQuest Table Name:	Libvirt.Virtual Disk
Open Table Name:	LIBVIRTVIRTUALDISK
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Errors	The number of I/O operation errors associated with the virtual disk [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Name	The name of the virtual disk device. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
ReadKB/s	The amount of data in kilobytes read by the virtual disk per second [Sequential = AVG Non-Sequential = SUM]
ReadRqst/s	The number of read requests per second associated with the virtual disk [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine (guest) to which the data applies. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
WriteKB/s	The amount of data in kilobytes written by the virtual disk per second [Sequential = AVG Non-Sequential = SUM]
WriteRqst/s	The number of write requests per second associated with the virtual disk [Sequential = AVG Non-Sequential = SUM]

7.12. Libvirt.Virtual Disk Config

The Libvirt.Virtual Disk Config table stores configuration information for virtual disk block devices in host systems.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Virtual Disk Config
IT Resource Name:	/TeamQuest/System/Open Virtualization/ <i>virtualmachinename</i>
TeamQuest Table Name:	Libvirt.Virtual Disk Config
Open Table Name:	LIBVIRTVDISKCONFIG
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
Allocation	The highest allocated extent of the virtual disk block device backing image in gigabytes [Non-Sequential = AVG]
Capacity	The logical size of the virtual disk block device backing image in gigabytes [Non-Sequential = AVG]
Name	The name of the virtual disk block device. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
Physical	The physical size of the container of the virtual disk block device backing image in gigabytes [Non-Sequential = AVG]

System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]

7.13. Libvirt.Virtual Machine

The Libvirt.Virtual Machine table stores configuration information for virtual machines associated with the collected host system.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Virtual Machine
IT Resource Name:	/TeamQuest/System/Open Virtualization/Virtual Machine/ <i>virtualmachinename</i>
TeamQuest Table Name:	Libvirt.Virtual Machine
Open Table Name:	LIBVIRTVM
Collection interval:	N/A
Default retentions:	1 month
Table type:	State

Statistic Name	Description
Autostart	The autostart state of the virtual machine. If the value is TRUE, the virtual machine is started when the host is started. [Non-Sequential = LST]
Id	The numeric identifier of the virtual machine. This identifier is assigned when the virtual machine is started. This field may report a value of -1 for inactive virtual machines. [Non-Sequential = ID]
MaxMemory	The maximum amount of memory the virtual machine is allowed to use in kilobytes [Non-Sequential = LST]
Memory	The memory used by the virtual machine in kilobytes [Non-Sequential = LST]
OSType	The operating system type of the virtual machine [Non-Sequential = ID]

Persistent	The persistence state of the virtual machine. If the value is TRUE, the virtual machine will continue to exist after it has been deactivated or if the host system is restarted. [Non-Sequential = LST]
State	<p>The current state of the virtual machine. This value can be any of the following:</p> <ul style="list-style-type: none">Blocked on a resourceCrashedNo StatePausedRunningShut offShutting downUnknown state <p>[Non-Sequential = ID]</p>
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
UUID	The universally unique identifier of the virtual machine [Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
VirtualCPUs	The number of virtual CPUs assigned to the virtual machine [Non-Sequential = SUM]

7.14. Libvirt.Virtual Machine Snapshot

The Libvirt.Virtual Machine Snapshot table stores configuration information for virtual machine snapshots.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Virtual Machine Snapshot
IT Resource Name:	/TeamQuest/System/Open Virtualization/Virtual Machine/ <i>virtualmachinename</i>
TeamQuest Table Name:	Libvirt.Virtual Machine Snapshot
Open Table Name:	LIBVIRTVMSNAPSHOT
Collection interval:	N/A
Default retentions:	1 month
Table type:	State

Statistic Name	Description
Name	The name of the virtual machine snapshot. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine (guest) to which the data applies. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]

7.15. Libvirt.Virtual Memory

The Libvirt.Virtual Memory table stores memory performance information for each virtual machine.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Virtual Memory
IT Resource Name:	/TeamQuest/System/Open Virtualization/Virtual Machine/ <i>virtualmachinename</i>
TeamQuest Table Name:	Libvirt.Virtual Memory
Open Table Name:	LIBVIRTVIRTUALMEMORY
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
ActualBalloon	The memory balloon value for the virtual machine in kilobytes. This field contains <N/A> if this functionality is not available. [Sequential = SUM Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Available	The total amount of usable memory by the virtual machine in kilobytes [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
MajorFault	The number of major page faults. A page fault occurs when a process accesses virtual memory that is not available. When servicing the page fault, if disk I/O operations are required, it is considered a major fault. [Sequential = SUM Non-Sequential = SUM]
MinorFault	The number of minor page faults. A page fault occurs when a process accesses virtual memory that is not available. When servicing the page fault, if disk I/O operations are not required, it is considered a minor fault. [Sequential = SUM Non-Sequential = SUM]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
SwapIn	The total amount of data read from swap space in kilobytes [Sequential = SUM Non-Sequential = SUM]
SwapOut	The total amount of data written to swap space in kilobytes [Sequential = SUM Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Unused	The amount of memory left unused by the virtual machine in kilobytes [Sequential = SUM Non-Sequential = SUM]
Virtual_Machine	The name of the virtual machine (guest) to which the data applies. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]

7.16. Libvirt.Virtual Network Config

The Libvirt.Virtual Network Config table stores configuration information for host virtual networks.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Virtual Network Config
IT Resource Name:	/TeamQuest/System/Open Virtualization/Virtual Machine/ <i>virtualmachinename</i>
TeamQuest Table Name:	Libvirt.Virtual Network Config
Open Table Name:	LIBVIRTVNETCONFIG
Collection interval:	N/A
Default retentions:	1 year
Table type:	State

Statistic Name	Description
Active	The state of the virtual network. This field contains True if the virtual network is active. [Non-Sequential = LST]
Autostart	The autostart state of the virtual network. This field contains True if the virtual network will be started when the virtual machine is started. [Non-Sequential = LST]
Bridge	The name of the network bridge associated with the virtual network [Non-Sequential = ID]
Name	The name of the virtual network. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Non-Sequential = ID]
Persistent	The persistence state of the virtual network. If the value is TRUE, the virtual network will continue to exist after it has been deactivated or if the host system is restarted. [Non-Sequential = LST]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
UUID	The universally unique identifier of the virtual network [Non-Sequential = ID]

7.17. Libvirt.Virtual Network Interface

The Libvirt.Virtual Network Interface table stores virtual machine performance data for virtual network interfaces.

Table Field Hierarchy

Class:	Libvirt
Subclass:	Virtual Network Interface
IT Resource Name:	/TeamQuest/System/Open Virtualization/Virtual Machine/ <i>virtualmachinename</i>
TeamQuest Table Name:	Libvirt.Virtual Network Interface
Open Table Name:	LIBVIRTVNETINTERFACE
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Name	The name of the virtual network interface. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
RcvDrops	The number of dropped receive packets for the virtual network interface device [Sequential = SUM Non-Sequential = SUM]
RcvErrors	The number of errors received by the virtual network interface device [Sequential = SUM Non-Sequential = SUM]
RcvKB/s	The amount of data received by the virtual network interface device in kilobytes [Sequential = SUM Non-Sequential = SUM]
RcvPackets/s	The number of packets received by the virtual network interface device [Sequential = SUM Non-Sequential = SUM]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies. This field is limited to 51 characters. Any value longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
XmitDrops	The number of dropped write packets for the virtual network interface device [Sequential = SUM Non-Sequential = SUM]
XmitErrors	The number of write errors for the virtual network interface device [Sequential = SUM Non-Sequential = SUM]
XmitKB/s	The amount of data writted by the virtual network interface device in kilobytes [Sequential = SUM Non-Sequential = SUM]
XmitPackets/s	The number of packets written by the virtual network interface device [Sequential = SUM Non-Sequential = SUM]

Section 8

Linux Systems

Notes:

- *Red Hat Linux is supported only for unmodified kernel releases that are part of a Red Hat Linux distribution.*
- *SuSE Linux is supported only for unmodified kernel releases that are part of a SuSE Linux distribution.*

Statistics for Linux systems are collected by the TeamQuest collection agents.

This section contains a listing of the statistics collected for the system:

- System Activity Statistics (see 8.1)
- Disk Space Statistics (see 8.2)
- Network Statistics (see 8.3)
- LPAR Configuration Statistics (see 8.4)
- Workload Statistics (see 8.5)
- Process Statistics (see 8.6)
- Hardware Inventory Statistics (see 8.7)
- System Log Statistics (see 8.8)
- General Log Statistics (see 8.9)
- TeamQuest Log Statistics (see 8.10)
- Derived Statistics (see 8.11)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV =Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

8.1. System Activity Statistics

The System Activity Agent is used to collect a wide variety of important system statistics. Major resources monitored by this agent include processors, memory, disks, network interfaces, and the operating system kernel.

Special Processing When Using Sequential Consolidation Method

Special processing occurs when certain records in the Block Device.by Device table are consolidated using the Sequential consolidation method. The following formulas are used to calculate the %busy, Actual_Interval, and record_count statistic values:

%busy

The %busy field uses a new consolidation method that uses the following formula to produce the consolidated %busy value:

$$\text{consolidated \%busy} = \%busy * \text{record_count} * \text{Actual_Interval}$$

At the end of the aggregation processing step after multiple records have been combined together to produce a single consolidated record, the %busy field contains the consolidated %busy value.

An additional processing step is performed using the following formula to produce a final %busy value that is stored into the consolidated record:

$$\%busy = \frac{\text{consolidated \%busy}}{\text{record_count} * \text{Interval}}$$

Note: The record_count field value used in the above formula must have already been generated using the record_count formula.

Block Device.by Device table records that have been stored by previous levels of TeamQuest collection agents do not contain the record_count field. For these records, a value of 1 is assumed for the record_count value.

Actual_Interval

For consolidated records (both reduced and not reduced), the Actual_Interval field should contain the Interval value at the end of the aggregation processing step.

record_count

The record_count field value is updated at the end of the aggregation processing step using the following formula:

$$\text{record_count} = \frac{(\sum(\text{record_count} * \text{Actual_Interval})) + (\text{Interval} - \sum \text{Actual_Interval})}{\text{Interval}}$$

The first part of the numerator is a summation of the record_count value multiplied by the Actual_Interval value across all of the consolidated records. This value should exist in the record_count field at the end of the aggregation processing step since the record_count field is weighted by Actual_Interval.

The Actual_Interval value used in the second part of the numerator is a summation of all the Actual_Interval values of the consolidated records. This value should exist in the Actual_Interval field at the end of the aggregation processing step since Actual_Interval is a summation.

The Interval value in the formula is the time range in seconds that the final consolidated record represents. For example, if 5-minute records are generated, the Interval value is calculated as 5 multiplied by 60.

Parameter Hierarchy

Class:	Block Device
Subclass:	by Device
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Block Device.by Device
Open Table Name:	BLKDEVBYDEVICE
Resource:	disk0, disk1, ...
Statistic Name:	
%busy	The percentage of time this device was servicing a transfer request [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-util.rpt
actq_avwait*	The average run queue wait time in milliseconds
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgresp*	The average response time of an I/O on a device. Calculated as await + aserv
avque	The average number of requests outstanding [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/io/dsk-q.rpt
aserv	The average time in milliseconds to service each transfer request (includes seek, rotation latency, and data transfer times) for the device [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-time.rpt
await	The average time in milliseconds that transfer requests are idle in the queue while the queue is occupied [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-time.rpt

Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
IO_intensity*	The activity of an I/O device. This is the product of the I/O response time in milliseconds and the I/O transfer rate in I/Os per second. This is proportional to the average queue length (the number of I/O requests waiting or in progress at the I/O device).
Kbytes/s	The rate at which data is transferred in kilobytes per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/hp-ux/sys-act/io/dsk-xfer.rpt /report/hp-ux/sys-act/io/top-dsk.rpt
record_count	The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record. [Sequential = AVG Non-Sequential = SUM]
reduction_name	The name of reduction rule [Sequential = ID Non-Sequential = ID]
reduction_source	The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H. [Sequential = ID Non-Sequential = ID]
transfers/s	The number of physical transfers to and from the disk per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/io/dsk-xfer.rpt
waitq_avwait*	The average wait queue wait time in milliseconds

Class:	Block Device
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Block Device.Summary
Open Table Name:	BLKDEVSUM
Statistic Name:	
transfers/s	The number of physical transfers to and from the disk per second [Sequential = AVG Non-Sequential = SUM]
Class:	CPU
Subclass:	by LPAR
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.by LPAR
Open Table Name:	CPUBYLPAR
Statistic Name:	
%entc	The percentage of the entitled processor capacity consumed [Sequential = AVG Non-Sequential = AVG]
%lpar_phys_busy	The percentage of the processor capacity consumed [Sequential = AVG Non-Sequential = SUM]
physc	The number of physical processors consumed [Sequential = AVG Non-Sequential = AVG]
Class:	CPU
Subclass:	by Processor
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.by Processor
Open Table Name:	CPUBYPROC
Resource:	cpu0, cpu1, ...
Statistic Name:	
%idle	The percentage of CPU time spent idle while no processes are waiting for I/O completion for the CPU [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/per-cpu.rpt
%guest	The percentage of CPU time spent running a virtual CPU for guest operating systems under the control of the Linux kernel. If the kernel level is below 2.6.32 or the system does not support virtualization, the value will be 0. [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/per-cpu.rpt

<code>%guest_nice</code>	<p>The percentage of CPU time spent running a virtual CPU for guest operating systems, in low priority, under the control of the Linux kernel. If the kernel level is below 2.6.33 or the system does not support virtualization, the value will be 0.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/linux/sys-act/cpu/per-cpu.rpt</p>
<code>%nice</code>	<p>The percentage of CPU time spent running low priority user processes for the CPU</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/linux/sys-act/cpu/per-cpu.rpt</p>
<code>%steal</code>	<p>The percentage of total CPU time that was utilized (stolen) by another virtual guest of the same real system</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/linux/sys-act/cpu/per-cpu.rpt</p>
<code>%sys</code>	<p>The percentage of CPU time spent running in system mode for the CPU</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/linux/sys-act/cpu/per-cpu.rpt</p>
<code>%usr</code>	<p>The percentage of CPU time spent running normal priority user processes for the CPU</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report: /report/linux/sys-act/cpu/per-cpu.rpt</p>
<code>%wio</code>	<p>The percentage of CPU time spent idle while some process is waiting for I/O completion</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>

Table Field Hierarchy

Class:	CPU
Subclass:	RelativePerformance
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.RelativePerformance
Open Table Name:	CPURELPERF
Collection interval:	1 minute
Default retentions:	1 month
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
rel_unused	The amount of CPU resources not used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
rel_used	The amount of CPU resources used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Class:	CPU
Subclass:	Summary
IT Resource Name:	/TeamQuest/Systethe partitionm/systemname/CPU
TeamQuest Table Name:	CPU.Summary
Open Table Name:	CPUSUM
Statistic Name:	
%busy	The percentage of time the CPU was not idle [Sequential = AVG Non-Sequential = AVG]
%guest	The percentage of total CPU time spent running a virtual CPU for guest operating systems under the control of the Linux kernel. If the kernel level is below 2.6.32 or the system does not support virtualization, the value will be 0. [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/cpu-util.rpt
%guest_nice	The percentage of total CPU time spent running a virtual CPU for guest operating systems, in low priority, under the control of the Linux kernel. If the kernel level is below 2.6.33 or the system does not support virtualization, the value will be 0. [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/cpu-util.rpt
%idle	The percentage of total CPU time spent idle while no processes are waiting for I/O completion [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/cpu-util.rpt
%nice	The percentage of total CPU time spent running low priority user processes [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/cpu-util.rpt
%steal	The percentage of total CPU time that was utilized (stolen) by another virtual guest of the same real system [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/cpu-util.rpt
%sys	The percentage of total CPU time spent in system mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/cpu-util.rpt
%usr	The percentage of total CPU time spent in user mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/cpu/cpu-util.rpt

%wio	The percentage of total CPU time spent idle while some process is waiting for I/O completion [Sequential = AVG Non-Sequential = AVG]
online_cpus	The number of CPUs that were online at the end of the sampling interval [Sequential = LST Non-Sequential = SUM]
Class:	Device Partition
Subclass:	by Partition
IT Resource Name:	/TeamQuest/System/systemname/Disk
TeamQuest Table Name:	Device Partition.by Partition
Open Table Name:	DEVPARTBYPARTITION
Resource:	partition0, partition1, ...
Statistic Name:	
%busy	The percentage of time the partition was servicing a transfer request. Available for kernel releases 2.6.25 and later. [Sequential = AVG Non-Sequential = AVG] View Reports: /report/linux/sys-act/io/par-util.rpt /report/linux/sys-act/io/topn-par.rpt
avque	The average number of requests outstanding per second. Available for kernel releases 2.6.25 and later. [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/io/par-q.rpt
avserv	The average service time in milliseconds for I/O requests for the partition. Available for kernel releases 2.6.25 and later. [Sequential = AVG Non-Sequential = AVG]
await	The average time in milliseconds for a transfer request to be completed by the partition. This includes time spent in queue and disk service time. Available for kernel releases 2.6.25 and later. [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/io/par-time.rpt
KB read/s	The number of kilobytes (KB) read from the partition per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/io/par-rdwr.rpt
KB write/s	The number of kilobytes (KB) transferred to the partition per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/io/par-rdwr.rpt

r+w/s	The number of read and write transfers to the partition per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/linux/sys-act/io/par-xfer.rpt /report/linux/sys-act/io/topn-par.rpt
reads/s	The number of read operations from the partition per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/io/par-rdwr.rpt
writes/s	The number of write operations to the partition per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/io/par-rdwr.rpt
Class:	Device Partition
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Device Partition.Summary
Open Table Name:	DEVPARTSUM
Statistic Name:	
transfers/s	The number of read and write transfers to all of the partitions per second [Sequential = AVG Non-Sequential = SUM]

Class: Kernel
Subclass: Load Average
IT Resource Name: /TeamQuest/System/*systemname*/Kernel
TeamQuest Table Name: Kernel.Load Average
Open Table Name: KNLLOADAVG
Statistic Name:
 1 min The number of processes in the run queue averaged over the last 1 minute
 [Sequential = LST Non-Sequential = SUM]
 View Report:
 /report/linux/sys-act/kernel/load-avg.rpt
 5 min The number of processes in the run queue averaged over the last 5 minutes
 [Sequential = LST Non-Sequential = SUM]
 View Report:
 /report/linux/sys-act/kernel/load-avg.rpt
 15 min The number of processes in the run queue averaged over the last 15 minutes
 [Sequential = LST Non-Sequential = SUM]
 View Report:
 /report/linux/sys-act/kernel/load-avg.rpt

Class: Kernel
Subclass: Paging
IT Resource Name: /TeamQuest/System/*systemname*/Memory
TeamQuest Table Name: Kernel.Paging
Open Table Name: KNLPAGING
Statistic Name:
 activepg The number of active (recently touched) pages in memory
 [Sequential = LST Non-Sequential = SUM]
 inactivepg The number of inactive pages in memory
 [Sequential = LST Non-Sequential = AVG]
 pgpgin/s The number of pages paged in per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/sys-act/kernel/paging.rpt
 pgpgout/s The number of pages paged out per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/sys-act/kernel/paging.rpt
 pgsize The size of the memory page
 [Sequential = LST Non-Sequential = SUM]

Class:	Kernel
Subclass:	Queues
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Queues
Open Table Name:	KNLQS
Statistic Name:	
runq_sz	The length of the run queue at the end of the sample interval [Sequential = AVG Non-Sequential = AVG]
cpuq_sz	The average length of the run queue per CPU at the end of the sample interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/sys-act/kernel/q-sizes.rpt
Class:	Kernel
Subclass:	Swapping
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Kernel.Swapping
Open Table Name:	KNLSWAPPING
Statistic Name:	
pgswpin/s	The number of pages swapped in per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/kernel/swapping.rpt
pgswapout/s	The number of pages swapped out per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/kernel/swapping.rpt
Class:	Kernel
Subclass:	System Call
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.System Call
Open Table Name:	KNLSYSCALL
Statistic Name:	
forks/s	The number of fork system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/syscall/imp-scal.rpt
pswch/s	The number of process switches per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/sys-act/syscall/imp-scal.rpt

Class:	Kernel
Subclass:	Tables
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Tables
Open Table Name:	KNLTABS
Statistic Name:	
proc-sz	<p>The number of entries presently used in the process table, which includes processes and threads. This count is taken at the end of the sampling interval.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/linux/sys-act/kernel/tbl-size.rpt</p>
Class:	Memory
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Memory
Open Table Name:	MEM
Statistic Name:	
buffermem	<p>The average amount of memory in use by buffers in megabytes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/linux/sys-act/memory/memory.rpt</p>
cachedmem	<p>The average amount of cached memory in use in megabytes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/linux/sys-act/memory/memory.rpt</p>
freemem	<p>The average amount of memory available to user processes in megabytes</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/linux/sys-act/memory/freemem.rpt</p> <p>/report/linux/sys-act/memory/memory.rpt</p>
highfree	<p>The amount of free memory in megabytes that is not directly mapped to kernel space</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
hightotal	<p>The total amount of memory in megabytes that is not directly mapped into kernel space</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
lowfree	<p>The amount of free memory in megabytes that is directly mapped into kernel space</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
lowtotal	<p>The total amount of memory in megabytes that is directly mapped into kernel space</p> <p>[Sequential = LST Non-Sequential = SUM]</p>

mlocked	The amount of memory in megabytes that is not available for paging. This statistic may be available on Linux systems with kernel releases 2.6.27 and later. [Sequential = LST Non-Sequential = SUM]
totalmem	The total amount of memory in megabytes [Sequential = LST Non-Sequential = SUM] View Report: /report/linux/sys-act/memory/memory.rpt
usedmem	The average amount of memory in use in megabytes [Sequential = LST Non-Sequential = SUM] View Report: /report/linux/sys-act/memory/memory.rpt
Class:	Swap Space
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/systemname/Memory
TeamQuest Table Name:	Swap Space.Summary
Open Table Name:	SWAPSPACESUM
Statistic Name:	
free	The amount of swap space in megabytes free at the end of the interval [Sequential = LST Non-Sequential = SUM] View Report: /report/linux/sys-act/swap/swpspc.rpt
in_use	The amount of swap space in megabytes in use at the end of the interval [Sequential = LST Non-Sequential = SUM]
total	The total amount of swap space in megabytes [Sequential = LST Non-Sequential = SUM] View Report: /report/linux/sys-act/swap/swpspc.rpt

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
bsp interval	The number of seconds elapsed between two data samples of the System Activity Agent [Sequential = SUM Non-Sequential = ID]
tqbsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqbsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

8.2. Disk Space Statistics

The Disk Space Agent tracks the disk space usage of locally mounted file systems. However, any data about Network File System (NFS) mounted file systems are not collected by the agent because the loss of the network connection to the NFS may result in hung processes.

Note: *File system names longer than 51 characters will be truncated.*

Parameter Hierarchy

Class:	Disk Space
Subclass:	by File System
IT Resource Name:	/TeamQuest/System/systemname/Disk
TeamQuest Table Name:	Disk Space.by File System
Open Table Name:	DISKSPACEBYFILESYS
Resource:	file system1, file system2, ...
Statistic Name:	
%inodes free*	The percentage of i-nodes available (not in use) on the file system at the end of the interval View Report: /report/linux/dskspace/total/pct-inod.rpt
%inodes used*	The percentage of i-nodes in use on the file system at the end of the interval View Report: /report/linux/dskspace/total/pct-inod.rpt
%space free*	The percentage of total space available (not in use) on the file system at the end of the interval View Reports: /report/linux/dskspace/io/total/pctspace.rpt /report/linux/dskspace/total/low-ones.rpt
%space used*	The percentage of total space in use on the file system at the end of the interval View Reports: /report/linux/dskspace/total/pctspace.rpt /report/linux/dskspace/total/fullest.rpt
%user space free*	The percentage of total user space available (not in use) on the file system at the end of the interval View Reports: /report/linux/dskspace/user/pctspace.rpt /report/linux/dskspace/user/low-ones.rpt
%user space used*	The percentage of total user space in use on the file system at the end of the interval View Report: /report/linux/dskspace/user/pctspace.rpt /report/linux/dskspace/user/fullest.rpt

capacity	The percentage of total space in use on the file system at the end of the interval [Sequential = LST Non-Sequential = AVG]
free (Mb)	The amount of space available (not in use) on the file system in megabytes (Mb) at the end of the interval including the space held back from normal users [Sequential = LST Non-Sequential = SUM] View Reports: /report/linux/dskspace/total/dskspace.rpt /report/linux/dskspace/total/low-ones.rpt
free inodes	The number of available (not in use) i-nodes on the file system at the end of the interval [Sequential = LST Non-Sequential = SUM] View Report: /report/linux/dskspace/total/i-nodes.rpt
total inodes	The total (used + available) number of i-nodes on the file system at the end of the interval [Sequential = LST Non-Sequential = SUM] View Report: /report/linux/dskspace/total/i-nodes.rpt
total (Mb)	The total (used + available) amount of space on the file system in megabytes (Mb) at the end of the interval including the space held back from normal users [Sequential = LST Non-Sequential = SUM] View Report: /report/linux/dskspace/total/dskspace.rpt
user free (Mb)	The amount of space available (not in use) on the file system in megabytes (Mb) at the end of the interval not including the space held back from normal users [Sequential = LST Non-Sequential = SUM] View Reports: /report/linux/dskspace/user/dskspace.rpt /report/linux/dskspace/user/low-ones.rpt
user total (Mb)*	The total (used + available) amount of space on the file system in megabytes (Mb) at the end of the interval not including the space held back from normal users View Report: /report/linux/dskspace/user/dskspace.rpt

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
dsp interval	The number of seconds elapsed between two data samples of the Disk Space Agent [Sequential = SUM Non-Sequential = ID]
tqdsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqdsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

8.3. Network Statistics

Network statistics are maintained in the TeamQuest performance database by the System Activity Agent. The statistics are classified by the hierarchy of key names.

Parameter Hierarchy

Class:	Network
Subclass:	by Interface
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	Network.by Interface
Open Table Name:	NETBYINTERFACE
Resource:	interface0, interface1, ...
Statistic Name:	
collisions/s	The number of network collisions per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-errs.rpt
in bytes/s	The number of network bytes input per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-byte.rpt
in drops/s	The number of network input drops per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-drop.rpt
in errors/s	The number of network input errors per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-errs.rpt
in packets/s	The number of network input packets per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-pkts.rpt
out bytes/s	The number of network bytes output per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-byte.rpt
out drops/s	The number of network output drops per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-drop.rpt

out errors/s	The number of network output errors per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-errs.rpt
out packets/s	The number of network output packets per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-pkts.rpt
Class:	Network
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/systemname/Network
TeamQuest Table Name:	Network.Summary
Open Table Name:	NETSUM
Statistic Name:	
errors/s	The total number of network errors per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]
in packets/s	The number of input packets per second summed from all network interfaces excluding the loop back interface [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-sum.rpt
out packets/s	The number of output packets per second summed from all network interfaces excluding the loop back interface [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/net-sum.rpt
total packets/s	The number of input and output packets per second summed from all network interfaces excluding the loop back interface [Sequential = AVG Non-Sequential = SUM]
Class:	NFSv2
Subclass:	Client
IT Resource Name:	/TeamQuest/System/systemname/Network
TeamQuest Table Name:	NFSv2.Client
Open Table Name:	NFSV2CLI
Statistic Name:	
calls/s*	The number of Network File System (NFS) version 2 calls per second sent by the client [Sequential = AVG Non-Sequential = SUM] View Report: /report/linux/network/clnt-v2.rpt

Linux Systems

Class: NFSv2
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv2.Client
Open Table Name: NFSV2CLI
Resource: create, getattr, link, lookup, mkdir, null, read, readdir, readlink, remove, rename, rmdir, root, statr, statfs, symlink, write, writecache
Statistic Name:
 reqs/s The number of NFS version 2 requests per second sent by the client
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/network/clnt-v2.rpt

Class: NFSv2
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv2.Server
Open Table Name: NFSV2SERV
Statistic Name:
 calls/s* The number of NFS version 2 calls per second received by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/network/srvr-v2.rpt

Class: NFSv2
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv2.Server
Open Table Name: NFSV2SERV
Resource: create, getattr, link, lookup, mkdir, null, read, readdir, readlink, remove, rename, rmdir, root, stat, statfs, symlink, write, writecache

Statistic Name:
reqs/s The number of NFS version 2 requests per second received by the server
[Sequential = AVG Non-Sequential = SUM]
View Report:
/report/linux/network/srvr-v2.rpt

Class: NFSv3
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Client
Open Table Name: NFSV3CLI

Statistic Name:
calls/s* The number of NFS version 3 calls per second received by the server
[Sequential = AVG Non-Sequential = SUM]
View Report:
/report/linux/network/srvr-v2.rpt

Class: NFSv3
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Client
Open Table Name: NFSV3CLI

Resource: access, commit, create, fsinfo, fsstat, getattr, link, lookup, mkdir, mknod, null, pathconf, read, readdir, readdir+, readlink, remove, rename, rmdir, setattr, symlink, write

Statistic Name:
reqs/s The number of NFS version 3 requests per second sent by the client
[Sequential = AVG Non-Sequential = SUM]
View Report:
/report/linux/network/clnt-v3.rpt

Class: NFSv3
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Server
Open Table Name: NFSV3SERV
Statistic Name:
 calls/s* The number of NFS version 3 calls per second received by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/network/srvr-v3.rpt

Class: NFSv3
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Server
Open Table Name: NFSV3SERV
Resource: access, commit, create, fsinfo, fsstat, getattr, link, lookup, mkdir, mknod, null, pathconf, read, readdir, readdir+, readlink, remove, rename, rmdir, setattr, symlink, write
Statistic Name:
 reqs/s The number of NFS version 3 requests per second received by the
 server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/network/srvr-v3.rpt

Class: RPC
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: RPC.Client
Open Table Name: RPCCLI
Statistic Name:
 authrefrsh/s The number of times authentication information had to be refreshed per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/network/rpc/client.rpt
 calls/s The number of RPC calls per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/network/rpc/client.rpt
 retrans/s The number of times a call had to be retransmitted due to a time-out while waiting for a reply from the server per second
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/linux/network/rpc/client.rpt

Class: RPC
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: RPC.Server
Open Table Name: RPCSERV
Statistic Name:
 badauth/s The total number of RPC calls per second
 [Sequential = AVG Non-Sequential = SUM]
 badcalls/s The number of calls rejected by the RPC layer per second. This is the sum of badlens/s and xdrcalls/s.
 [Sequential = AVG Non-Sequential = SUM]
 badclnt/s The number of calls per second rejected by the RPC layer because of a bad client
 [Sequential = AVG Non-Sequential = SUM]
 calls/s The number of RPC calls received per second
 [Sequential = AVG Non-Sequential = SUM]
 xdrcalls/s The number of RPC calls per second whose header could not be External Data Representation (XDR) decoded
 [Sequential = AVG Non-Sequential = SUM]

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
tqwarp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqwarp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]
warp interval	The number of seconds elapsed between two data samples of the Process-Workload Agent [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

8.4. LPAR Configuration Statistics

The following statistics are stored in the TeamQuest performance database tables by the System Activity Agent.

Note: *LPAR statistics are supported only on Linux POWER systems.*

Table Field Hierarchy

Class:	IBM POWER
Subclass:	LPAR
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	IBM POWER.LPAR
Open Table Name:	IBMPOWERLPAR
Collection interval:	N/A
Default retention:	1 year
Table type:	Event

Statistic Name	Description
active_cpus_in_pool	The current number of active physical CPUs in the shared processor pool being used by this LPAR [Non-Sequential = SUM]
active_physical_cpus_in_system	The current number of active physical CPUs in the system containing this LPAR [Non-Sequential = SUM]
capacity_increment	The granule at which changes to the Entitled Capacity can be made. A value in whole multiples indicates a dedicated LPAR. [Non-Sequential = NON]
entitled_capacity	The number of processing units this LPAR is entitled to receive [Non-Sequential = SUM]
maximum_capacity	The maximum number of processing units this LPAR was defined to ever have. Entitled capacity can be increased up to this value. [Non-Sequential = SUM]
maximum_memory	The maximum possible amount of memory [Non-Sequential = SUM]
maximum_physical_cpus_in_system	The maximum possible number of physical CPUs in the system containing this LPAR [Non-Sequential = MAX]
maximum_virtual_cpus	The maximum possible number of CPUs (virtual engines). This statistic is not available from the Linux operating system and is set to <N/A>. [Non-Sequential = SUM]
minimum_capacity	The minimum number of processing units this LPAR was defined to ever have. Entitled capacity can be reduced down to this value. [Non-Sequential = SUM]

minimum_memory	The minimum amount of memory this LPAR was defined to ever have. This statistic is not available from the Linux operating system and is set to <N/A>. [Non-Sequential = SUM]
minimum_virtual_cpus	The minimum possible number of CPUs this LPAR was defined to ever have [Non-Sequential = SUM]
mode	Indicates whether the LPAR processor capacity is capped, or uncapped and allowed to consume idle cycles from the shared pool. A dedicated LPAR is implicitly capped. [Non-Sequential = ID]
os_name	The system name as assigned by the operating system [Non-Sequential = ID]
online_memory	The amount of memory currently online [Non-Sequential = SUM]
online_virtual_cpus	The number of CPUs (virtual engines) currently online [Non-Sequential = SUM]
partition_group_id	The identifier of the shared pool of physical processors of which the LPAR is a member [Non-Sequential = ID]
partition_name	The logical partition name as assigned by the hardware management console (HMC) [Non-Sequential = ID]
partition_number	The number of the logical partition [Non-Sequential = ID]
physical_cpu_percentage	The fractional representation relative to whole physical CPUs that the LPAR virtual CPUs equate to. This is a function of Entitled Capacity/Online CPUs. Dedicated LPARs would have a Physical CPU Percentage of 100%. A 4-way virtual with Entitled Capacity of 2 processor units would have a Physical CPU Percentage of 50%. [Non-Sequential = SUM]
shared_pool_id	The LPAR group of which the LPAR is a member of which the LPAR is a member. This statistic is not available from the Linux operating system and is set to <N/A>. [Non-Sequential = ID]
smt	Indicates whether simultaneous multi-threading (SMT) is enabled. This statistic is not available from the Linux operating system and is set to <N/A>. [Non-Sequential = ID]
system_name	The name of the physical system [Non-Sequential = ID]
timestamp	The timestamp when the data was collected [Non-Sequential = LST]
type	Indicates whether the LPAR is using dedicated or shared CPU resources [Non-Sequential = ID]

unallocated_capacity	The number of processor units currently unallocated in the shared processor pool being used by this LPAR [Non-Sequential = NON]
variable_capacity_weight	The priority weight assigned to this LPAR, which controls how extra (idle) capacity is allocated to it. A weight of -1 indicates a soft cap is in place. [Non-Sequential = NON]

8.5. Workload Statistics

Workload statistics are maintained in the TeamQuest performance database by the Process-Workload Agent. The statistics are classified by the hierarchy of key names.

Parameter Hierarchy

Class:	Workload
Subclass:	by Workload
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /workload/ <i>workloadset</i> / <i>workload</i>
TeamQuest Table Name:	Workload.by Workload
Open Table Name:	WLBYWORKLOAD
Workload Set:	WLS1, WLS2, ...
Workload:	WL1, WL2, ...
Statistic Name:	
%cpu	The percentage of total CPU consumed by the workload. Total CPU time is the value of the sampling interval multiplied by the number of CPUs on the system. Thus, if the sum of %cpu for all workloads is less than 100%, some CPUs must have been idle for some time during the sampling interval. [Sequential = AVG Non-Sequential = AVG] View Report: /report/linux/workload/pct-cpu.rpt
avgmem	The cumulative swap process image size in kilobytes of the running processes in the workload at the end of the sampling interval [Sequential = AVG Non-Sequential = SUM]
etime	The sum of the elapsed times in seconds of all of the processes in the workload. Dividing this number by the number of processes in the workload (ponging + pcomplete) gives the average time a process in the workload existed during the sampling interval. [Sequential = SUM Non-Sequential = SUM]
liocho	The number of logical characters in kilobytes transferred by the workload during the sampling interval. Available for kernel releases 2.6.20 and later. [Sequential = SUM Non-Sequential = SUM]

majflt	<p>The number of major page faults generated by the workload for processes that were active at the end of the sampling interval. A major page fault is a page fault that requires I/O.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/maj-flt.rpt</p>
pcomplete	<p>The number of processes completed in the sampling interval. For process data, the same number is called cproc.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/num-proc.rpt</p>
pio	<p>The number of physical I/O transfers done by the workload during the sampling interval. The number reported represents only the completed processes during the sampling interval. This value is always zero for Linux systems.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/pio.rpt</p>
pongoing	<p>The number of processes running at the end of the sampling interval. In process data, the same number is called oproc. The number of processes in a workload could be derived by adding pongoing and pcomplete. This sum is called nproc in process data.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/num-proc.rpt</p>
prss	<p>The resident set size in kilobytes of private memory occupied by the running processes in the workload at the end of the sampling interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/rss.rpt</p>
pstart	<p>The number of processes started in the sampling interval. In process data, this number is called sproc.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/num-proc.rpt</p>
record_count	<p>The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
reduction_name	<p>The name of reduction rule</p> <p>[Sequential = ID Non-Sequential = ID]</p>
reduction_source	<p>The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H.</p> <p>[Sequential = ID Non-Sequential = ID]</p>

rss	<p>The resident set size (real memory) in kilobytes of all processes running at the end of the interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/rss.rpt</p>
rwKB	<p>The number of kilobytes that the workload read from and wrote to disk devices per sample. The kernel may read and write more data than requested by the workload. This value will represent the amount that the kernel read and wrote. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
rwKBt	<p>The number of kilobytes the workload read from and wrote to page cache, disks, and terminals per sample. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
rwsysc	<p>The number of read and write system calls made by the workload per sample. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
srss	<p>The resident set size in kilobytes of shared memory occupied by the running processes in the workload at the end of the sampling interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/rss.rpt</p>
syscpu	<p>The system CPU time in seconds used by the workload. System CPU time is the time spent in kernel mode (for example, the time spent in executing system calls, paging, and so on).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/sys-cpu.rpt</p>
threads	<p>The number of threads at the end of the sampling interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/threads.rpt</p>
totcpu	<p>The total CPU time in seconds used by the workload. This value is the same as the sum of usrcpu + syscpu.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/cpu-util.rpt</p>
usrcpu	<p>The user CPU time in seconds used by the workload. User CPU time is the time the CPU spent running in user mode.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/linux/workload/user-cpu.rpt</p>

8.6. Process Statistics

The Process-Workload Agent collects process data from the operating system and processes accounting files. The Process-Workload Agent calculates the usage of every process in a given interval, applies reduction definitions to each process, and stores the reduced process data. It also applies workload definitions to the reduced process data and stores system resource usage by workload.

Reduced Process Records

The Process-Workload Agent collects process data and reduces the data according to the user-defined reduction definitions. A reduction definition may cause multiple processes to be merged into a single process record. Thus, a process record contains data about one or more processes. When you are looking at the resource usage numbers, it is important to know how many processes a process record actually represents. The `nproc` data item indicates exactly how many processes each process record is representing. When a process record is representing more than one process, the resource usage fields such as `totcpu`, `rss`, and `pio_t` are the sum of the resource usage of the individual processes. When all of the processes do not have the same value for a field, the identifier fields such as `command`, `login`, and `gid` are set to `<Multi>`. When data for some fields is not available, the fields are set to `<N/A>`.

Disabling Reduction Definitions

If you want to look at the details of every individual process and do not wish to have merged process records, you must disable reduction processing by making all reduction sets inactive. However, with reduction processing disabled, more records have to be stored and more disk space is needed. For information on disabling reductions, see the *TeamQuest Performance Software Administration Guide*.

Process Data with Process Accounting Turned Off

If process accounting is turned off, the process data is incomplete as data about completed processes is not available. In this case, the process data only shows a portion of the activity in the sampling interval. To find out whether process accounting is turned on or off, look at the `cproc` field of all of the process records. If the `cproc` field indicates all zeros, it means that processes were not completed in that sampling interval and that the process accounting is turned off.

When process accounting is not turned on, the process record, `<Other>` includes the CPU time for processes that completed during the interval.

For more information on process accounting, see the *TeamQuest Performance Software Administration Guide*.

Retrieving Hardware Configuration Information

The Process-Workload Agent retrieves hardware configuration information. The information is stored upon startup and once-a-day in the `HINV.Summary`, `HINV.Devices`, and `HINV.FileSystem` table files of the TeamQuest performance database. The information is also stored when starting the agent and when the agent detects a change in configuration.

Retrieving Process I/O Statistics

When file access to I/O statistics (/proc/<pid>/io) is restricted by the Linux kernel, non-root users cannot collect process I/O statistics for other users. These statistics are stored as <N/A>. However, running the Process-Workload Agent as root collects process I/O statistics for all users.

Table Field Hierarchy

Class:	LINUX
Subclass:	Process
IT Resource Name:	/TeamQuest/System/systemname/Process
TeamQuest Table Name:	LINUX.Process
Open Table Name:	LNXPROC
Collection interval:	Based on the primary aggregation set
Default retention:	1 day
Table type:	Performance

Note: The collection interval is also dependent on the Processes Only setting in the configuration file for the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgmem_t	The size of the swappable process image in kilobytes. If a process starts and ends within the same interval, the number is unavailable and is marked as <N/A>. [Sequential = AVG Non-Sequential = SUM]
btime	The start time of the process. For process records representing more than one process, this field shows the earliest of the start times. [Sequential = FST Non-Sequential = ID]
command	The command name of the process. If a process starts and ends within the same interval, only up to 8 characters of the command name can be displayed. Otherwise, up to 32 characters are displayed. Therefore, an “automountd” process may appear as “automoun” if it starts and ends within the same interval. In such cases, you may want to set up your workload, reduction, or filter definitions to catch both the command names as in the following: command = {“automoun”, “automountd”}. [Sequential = ID Non-Sequential = ID]
cproc	The number of processes completed in the interval [Sequential = SUM Non-Sequential = SUM]

cwriteKB	<p>The number of kilobytes that the process caused to not happen by truncating page cache. For example, if the process writes 1 MB to a file and deletes the file, the write function to the disk is canceled because the file has been deleted. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
cwriteKB_t	<p>The number of kilobytes that the process caused to not happen by truncating page cache since the process started. For example, if the process writes 1 MB to a file and deletes the file, the write function to the disk is canceled because the file has been deleted. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
etime	<p>The elapsed time in seconds for the current interval. This number tells how long a process existed in the current interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
etime_t	<p>The total elapsed time in seconds. This number tells how long a process existed since it started. For a single process, this is the same as the value of the Time field minus the value of the btime field.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
fullecmd	<p>The full command string, including arguments, for the process. If a process begins and ends in the sampling interval, the field is unavailable and is marked as <N/A>. This is an important consideration when using a workload, reduction, or filter definition. The number of characters that are stored is determined by the Command Length setting of the Process-Workload Agent configuration file in TeamQuest Manager. You can also have either the first or the last <i>N</i> characters of the command displayed. The setting is specified by the Command Orientation setting of the Process-Workload Agent configuration file in TeamQuest Manager. The limit of characters from the operating system data source is 4095. For a description of how this may affect collection when the LAST Command Orientation is configured for the Process-Workload Agent, see the <i>TeamQuest Performance Software Administration Guide</i>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
gid	<p>The real group identifier of the owner of the process</p> <p>[Sequential = ID Non-Sequential = ID]</p>
group	<p>The group name of the owner of the process. This field is derived from gid.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Interval	<p>The expected data sampling interval</p> <p>[Sequential = SUM Non-Sequential = ID]</p>
login	<p>The login name of the owner of the process. This field is derived from uid.</p> <p>[Sequential = ID Non-Sequential = ID]</p>

majflt	<p>The number of major page faults generated in the current interval. A major page fault is a page fault that requires I/O. If a process ends in the sampling interval, the number is unavailable and is marked as <N/A>.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
majflt_t	<p>The total number of major page faults generated since the process started. A major page fault is a page fault that requires I/O.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
nproc	<p>The number of processes that the process record represents. If a process starts and ends with the same interval, the number is unavailable and is marked as <N/A>. In a reduced record, it is the number of processes that were merged together to form a single process record. When time consolidation is not applied to the process data, this number should equal the sum of cproc and oproc. When consolidating over time, the number represents the average number of processes in the process record for the consolidated period. An average is used because processes can move in and out and between reduced records from interval to interval.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
oproc	<p>The number of ongoing processes at the end of the interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
pctcpu	<p>The percentage of total available CPU time the process used in the current sampling interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
pid	<p>The process identifier number. If a process starts and ends within an interval, this number is unavailable and is marked as <N/A>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
pio	<p>The number of physical I/O requests for the current interval. The number reported represents only the completed processes during the sampling interval. This value is always zero for Linux systems.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
pio_t	<p>The total number of physical I/O requests since the process started. The number reported represents only the completed processes during the sampling interval.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
ppid	<p>The numerical identifier of the parent of a process. If a process starts and ends within an interval, the number is unavailable and is marked as <N/A>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
pri	<p>The priority of the process. Higher numbers mean lower priority. If a process ends within an interval, the number is unavailable and is marked as <N/A>. If the process record represents more than one process, the priority is an average of the constituent running processes' priority values.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>

prss	The resident set size in kilobytes of private memory for the process at the end of the sampling interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>. [Sequential = AVG Non-Sequential = SUM]
readKB	The number of kilobytes that the process read from disk devices per sample. The kernel may read more data than requested by the process. This value will represent the amount that the kernel read. Available for kernel releases 2.6.20 and later. [Sequential = SUM Non-Sequential = SUM]
readKB_t	The number of kilobytes that the process read from disk devices since the process started. The kernel may read more data than requested by the process. This value will represent the amount that the kernel read. Available for kernel releases 2.6.20 and later. [Sequential = LST Non-Sequential = SUM]
readKBt	The number of kilobytes the process read from page cache, disks, and terminals per sample. Available for kernel releases 2.6.20 and later. [Sequential = SUM Non-Sequential = SUM]
readKBt_t	The number of kilobytes the process read from page cache, disks, and terminals since the process started. Available for kernel releases 2.6.20 and later. [Sequential = LST Non-Sequential = SUM]
redname	The reduction name of the process record. If a process did not match any of the reduction definitions, then it would not be reduced and the field will be blankthe running. [Sequential = ID Non-Sequential = ID]
rss	The resident set size (real memory size) of the running process at the end of the interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>. [Sequential = AVG Non-Sequential = SUM]
rsysc	The number of read system calls made by the process per sample. Available for kernel releases 2.6.20 and later. [Sequential = SUM Non-Sequential = SUM]
rsysc_t	The number of read system calls made by the process since the process started. Available for kernel releases 2.6.20 and later. [Sequential = LST Non-Sequential = SUM]
rwKB	The number of kilobytes that the process read from and wrote to disk devices per sample. The kernel may read and write more data than requested by the process. This value will represent the amount that the kernel read and wrote. Available for kernel releases 2.6.20 and later. [Sequential = SUM Non-Sequential = SUM]
rwKB_t	The number of kilobytes that the process read from and wrote to disk devices per sample since the process started. The kernel may read and write more data than requested by the process. This value will represent the amount that the kernel read and wrote. Available for kernel releases 2.6.20 and later. [Sequential = LST Non-Sequential = SUM]

rwKBt	<p>The number of kilobytes the process read from and wrote to page cache, disks, and terminals. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
rwKBt_t	<p>The number of kilobytes that the process read from and wrote to the page cache, disks, and terminals since the process started. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
rwsysc	<p>The number of read and write system calls made by the process per sample. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
rwsysc_t	<p>The number of read and write system calls made by the process since the process started. Available for kernel releases 2.6.20 and later.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
sproc	<p>The number of processes started in the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
srss	<p>The resident set size in kilobytes of shared memory occupied by the running processes at the end of the sampling interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
syscpu	<p>The system CPU time in seconds for the current interval. System CPU time is the time the CPU spent running in kernel mode (for example, the time spent in executing system calls, paging, and so on). If an application is taking a lot of syscpu time, you may want to optimize the use of system calls (for example, use a larger block size for I/O).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
syscpu_t	<p>The total system CPU time in seconds</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
threads	<p>The number of threads associated with the process at the end of the interval. This value represents the number of light-weight processes (LWP) or kernel-supported user threads. A thread is a dynamic object that represents a control point in a process and executes a sequence of instructions.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>

totcpu	The total CPU time in seconds used in the current interval. This number is the same as the sum of usrcpu and syscpu. [Sequential = SUM Non-Sequential = SUM]
totcpu_t	The total CPU time (user + system) in seconds used by the process since it started [Sequential = LST Non-Sequential = SUM]
tty	The controlling terminal identifier in dev_t format. For the processes without a controlling terminal, this field will contain a -1. [Sequential = ID Non-Sequential = ID]
ttyname	The controlling terminal for the process. It is a device name without the /dev/ prefix. This is derived from tty. For the processes without a controlling terminal, this field will contain a question mark (?). [Sequential = ID Non-Sequential = ID]
uid	The real user id of the process owner [Sequential = ID Non-Sequential = ID]
usrcpu	The user CPU time in seconds for the current interval. User CPU time is the time the CPU spent running in user mode. If an application is taking a lot of usrcpu time, you should try to optimize the code, if possible. [Sequential = SUM Non-Sequential = SUM]
usrcpu_t	The total user CPU time in seconds since the start of the process [Sequential = LST Non-Sequential = SUM]
Workload	<p>The workload set and the workload associated with the process. When the Process-Workload Agent stores the process record, this field is blank. When process records are reported, the workload can be evaluated and is shown in the report.</p> <p>This field is available for reporting only when using TeamQuest Analyzer and TeamQuest tView.</p> <p>Workload evaluation takes place when data is retrieved from the database, based on workload sets defined in the database where the data is stored. Workload sets reported in the Workload field do not have to be active.</p> <p>For more information on workload evaluation, see the <i>TeamQuest Analyzer User Guide</i> or the <i>TeamQuest Performance Software Command Line Interfaces Reference Manual</i>. [Sequential = ID Non-Sequential = ID]</p>
workload:wlsname	<p>There is one field for each <i>wlsname</i> (Workload Set Name). The value for this field shows the name of the workload to which the process belongs. If a process belongs to none of the workloads, it will display the workload name "OTHER."</p> <p>This field is available for reporting only when using TeamQuest View or TeamQuest cView. [Sequential = ID Non-Sequential = ID]</p>

writeKB	The number of kilobytes that the process wrote to disk devices per sample. The kernel may write more data than requested by the process. This value will represent the amount that the kernel wrote. Available for kernel releases 2.6.20 and later. [Sequential = SUM Non-Sequential = SUM]
writeKB_t	The number of kilobytes that the process wrote to disk devices since the process started. The kernel may write more data than requested by the process. This value will represent the amount that the kernel wrote. Available for kernel releases 2.6.20 and later. [Sequential = LST Non-Sequential = SUM]
writeKBt	The number of kilobytes written by the process to page cache, disks, and terminals per sample. Available for kernel releases 2.6.20 and later. [Sequential = SUM Non-Sequential = SUM]
writeKBt_t	The number of kilobytes the process wrote to page cache, disks, and terminals since the process started. Available for kernel releases 2.6.20 and later. [Sequential = LST Non-Sequential = SUM]
wsync	The number of write system calls made by the process per sample. Available for kernel releases 2.6.20 and later. [Sequential = SUM Non-Sequential = SUM]
wsync_t	The number of write system calls made by the process since the process started. Available for kernel releases 2.6.20 and later. [Sequential = LST Non-Sequential = SUM]

8.7. Hardware Inventory Statistics

The Process-Workload Agent retrieves hardware configuration information. The information is stored upon startup and once-a-day in the HINV.Summary, HINV.CPUModel, HINV.CPU Thread Speeds, HINV.Devices, and HINV.FileSystem table files of the TeamQuest performance database. The information is also stored if the agent detects a change in configuration.

Note: *The storage of hardware inventory records depends on the Hardware Inventory setting in the configuration file of the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.*

Table Field Hierarchy

Class:	HINV
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	HINV.Summary
Open Table Name:	HINVSUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
core_multi_thread	The status or ability of the processor to support multiple independent threads. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cores_per_chip	The number of cores or processors on an individual chip. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cpu_chips	The number of CPU chips or sockets. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cpu_count	The number of configured processors. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cpu_speed	The speed of the processor in MHz or GHz [Non-Sequential = ID]
cpu_type	The basic instruction set architecture of the current system [Non-Sequential = ID]
logical_cpu_count	The number of logical processors [Non-Sequential = ID]

mem_size	The size of configured random access memory in kilobytes where 1 kilobyte = 1,024 bytes [Non-Sequential = ID]
memory	The size of configured random access memory in megabytes where 1 megabyte = 1,048,576 bytes [Non-Sequential = ID]
memory_size	The size of configured random access memory in megabytes or gigabytes [Non-Sequential = ID]
model	Name of the hardware implementation or platform [Non-Sequential = ID]
os_release	The name and level of this implementation of the operating system [Non-Sequential = ID]
pagesize	The size of a page of memory [Non-Sequential = ID]
partition_type	The partition type of the system. The value indicates the system hypervisor type, guest type, logical partition type, zone type, or logical domain type. If the system does not have a partition type, this field will be blank. [Non-Sequential = ID]
serial	The hardware-specific serial number of the physical machine [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_identifier	Information used to identify the system [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
timezone	The time zone where the data was collected [Non-Sequential = ID]
TQLevel	The level of TeamQuest Manager [Non-Sequential = ID]

The HINV.CPUModel table stores best-match, relative performance data about the system configuration. This table is created by the Hardware Inventory Agent (**tqhin**v) to map physical hardware to a CPU model that describes performance in relative terms. This table is not created for any virtualized system. It is populated for physical systems only. It is not populated for VMware guests, Hyper-V guests, Solaris logical domains (LDMs), Solaris guest LDMs, KVM guests, and Linux on POWER systems.

Table Field Hierarchy

Class:	HINV
Subclass:	CPUModel
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.CPUModel
Open Table Name:	HINVCPUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
cpu_chips	The number of CPU chips or sockets [Non-Sequential = ID]
cpu_confidence	The percentage of confidence in the correctness of the CPU match based on model, frequency, and configuration (chips, cores, threads) [Non-Sequential = SUM]
cpu_cores	The number of CPU cores or processors on an individual CPU chip [Non-Sequential = ID]
cpu_name	The name of the selected CPU [Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Non-Sequential = ID]
cpu_speed	The speed of the processor in megahertz (MHz) or gigahertz (GHz) [Non-Sequential = ID]
cpu_threads	The number of CPU threads on an individual CPU core or processor [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
user_override	The user override status of the default TeamQuest generated CPU match. This field is not currently used and should appear as 0. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	CPU Thread Speeds
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.CPU Thread Speeds
Open Table Name:	HINVCPUTHREADSPEEDS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
speed_up_factor	The performance improvement when there are multiple active threads per core, compared to when there is only one active thread per core [Sequential = AVG Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
thread_number	The number of active threads [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = ID Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	Devices
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.Devices
Open Table Name:	HINVDEVS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
class	The device classification: controller, disk, or tape [Non-Sequential = ID]
controller	The device path indicator which defines a connection to another device [Non-Sequential = ID]
lun_id	The globally unique Logical Unit Number (LUN) identifier for Storage Area Network (SAN) based disk devices. This field is blank for non-SAN based disk devices, CD-ROM drives, tape drives, and so on. [Non-Sequential = ID]
name	The unique identifier for the device [Non-Sequential = ID]
name2	The alternate device name. This field may be blank. [Non-Sequential = ID].
product	The product identifier. This field may be blank. [Non-Sequential = ID]
revision	The revision level for this product. This field may be blank. [Non-Sequential = ID]
rpm	The speed at which the media spins. If an actual value cannot be obtained for the device, a default value of 7,200 is used. [Non-Sequential = ID]
sequence	The sequence number of the device [Non-Sequential = ID]
swap	A true or false statement which indicates whether or not a swap file exists on the device [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp when the data was collected [Non-Sequential = ID]
vendor	The name of the device vendor. This field may be blank. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	FileSystem
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.FileSystem
Open Table Name:	HINVFILESYS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
BlkSize	The size of a block on the file system [Non-Sequential = ID]
Device	The path for the device on which the file system is mounted [Non-Sequential = ID]
Name	The unique identifier for the file system [Non-Sequential = ID]
Source	The source physical disk or logical volume of the file system. This field is always blank for this platform. [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
TotBlks	The total number of blocks on the file system [Non-Sequential = ID]
TotFiles	The maximum total number of files, as represented by inodes, possible on the file system. Some inodes may be used for entities other than visible files. [Non-Sequential = ID]
TotSize	The total amount of space on the file system in megabytes [Non-Sequential = ID]
Type	The type of the file system [Non-Sequential = ID]

8.8. System Log Statistics

The System Log Agent is used to collect system log messages generated by the system log daemon (syslogd). The System Log Agent stores these messages in the TeamQuest performance database for analysis and alarm reporting. The log messages are separated into four fields; the time that the message was posted, the host system from where the message was initiated, the program or user that posted the message, and the text of the message.

Table Field Hierarchy

Class:	System
Subclass:	System Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /System Log
TeamQuest Table Name:	System.System Log
Open Table Name:	SYSSYSTEMLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Event_Time	The time that the message was logged to the system log [Non-Sequential = ID]
Loghost	The name of the system that logged the message [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Reporter	The name of the user or process that logged the message [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system where the log message originated. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

8.9. General Log Statistics

The General Log Agent is used to collect log messages generated by application programs. The General Log Agent stores these messages in the TeamQuest performance database for analysis. Examples include backup, security, database, and Web server applications.

Table Field Hierarchy

Class:	System
Subclass:	General Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /General Log
TeamQuest Table Name:	System.General Log
Open Table Name:	SYSGENERALLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The message type [Non-Sequential = ID]

8.10. TeamQuest Log Statistics

The following statistics are stored in the performance database tables by the TeamQuest Log Agent. The collection interval and retention periods can be modified. For more information on modifying the collection interval and retention periods, see the *TeamQuest Performance Software Administration Guide*.

Table Field Hierarchy

Class:	Service
Subclass:	TeamQuest Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /TeamQuest Log
TeamQuest Table Name:	Service.TeamQuest Log
Open Table Name:	SVCTQLOG
Collection interval:	N/A
Default retention:	1 day
Table type:	Event

Statistic Name	Description
Filename	The name of the TeamQuest log file that was the source of the message text [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The log message type. This is always set to tqlog . [Non-Sequential = ID]

8.11. Derived Statistics

Some products within TeamQuest Performance Software use derived statistics to display common statistics across different platforms. The derived statistics are inserted into the performance database when databases are created. In this section, a derived statistic is marked with an asterisk (*).

You can find information on the following derived statistics:

- Workload Performance Derived Statistics
- TeamQuest On the Web Derived Statistics
- TeamQuest Alert Derived Statistics

8.11.1. Workload Performance Derived Statistics

TeamQuest Manager maintains the following derived statistics that use data from the System Activity Agent and the Process-Workload Agent. The workload performance reports reference these statistics. For information on workload performance reports, see the *TeamQuest View Reports Reference Manual*.

Parameter Hierarchy

Class:	Derived
Subclass:	Workload Performance.by Workload
Workload Set:	WLS1, WLS2, ...
Workload:	ALL
Statistic Name:	
%cpu*	The total percentage of CPU utilization. Collected by the Process-Workload Agent.
Kbytes resident memory/process*	The average amount of resident memory used per process. Collected by the Process-Workload Agent.
Kbytes virtual memory/process*	The average amount of virtual memory used per process. Collected by the Process-Workload Agent.
PIOs/sec*	The number of physical I/Os per second. Collected by the Process-Workload Agent.
Population (etime/interval)*	The average number of concurrent processes. Collected by the Process-Workload Agent.
Response (etime/process)*	The elapsed time per process. Collected by the Process-Workload Agent.
Throughput (processes/sec)*	The number of processes completed per second. Collected by the Process-Workload Agent.
Total Kbytes resident memory*	The average amount of resident memory used by the workload. Collected by the Process-Workload Agent.
Total Kbytes virtual memory*	The average amount of virtual memory used by the workload. Collected by the Process-Workload Agent.

Class:	Derived
Subclass:	Workload Performance.Summary
Workload Set:	WLS1, WLS2, ...
Statistic Name:	
%cpu*	The total percentage of CPU utilization. Collected by the Process-Workload Agent.
Kbytes resident memory/process*	The average amount of resident memory used per process. Collected by the Process-Workload Agent.
Kbytes virtual memory/process*	The average amount of virtual memory used per process. Collected by the Process-Workload Agent.
PIOs/sec*	The number of physical I/Os per second. Collected by the Process-Workload Agent.
Population (etime/interval)*	The average number of concurrent processes. Collected by the Process-Workload Agent.
Response (etime/process)*	The elapsed time per process. Collected by the Process-Workload Agent.
Throughput (processes/sec)*	The number of processes completed per second. Collected by the Process-Workload Agent.
Total Kbytes resident memory*	The average amount of resident memory used by the workload. Collected by the Process-Workload Agent.
Total Kbytes virtual memory*	The average amount of virtual memory used by the workload. Collected by the Process-Workload Agent.

8.11.2. TeamQuest On the Web Derived Statistics

The derived statistics used by TeamQuest On the Web include the following:

Parameter Hierarchy

Class: Derived

Subclass: TQWeb.Summary

Statistic Name:

disk_xfers_per_sec*	The total number of read and write transfers per second for all devices
free_disk_space*	The amount of space available (not in use) on all file systems in megabytes. This measurement is taken at the end of the sampling interval and includes the space held back from normal users.
free_real_mem*	The amount of free memory available in megabytes. This measurement is taken at the end of the sampling interval.
free_swap_space*	The number of megabytes free for process swapping
pct_cpu_busy*	The percentage of total CPU time the CPU was busy (not idle). This value includes the time running system code and the time running normal priority user processes.
pct_sys_cpu*	The percentage of total CPU time spent in system mode
pct_usr_cpu*	The percentage of total CPU time spent running in user mode
pkt_errors_per_sec*	The total number (in + out) of network errors per second for all network interfaces
pkts_in_per_sec*	The total number of network input packets per second for all network interfaces
pkts_out_per_sec*	The total number of network output packets per second for all network interfaces
pkts_per_sec*	The total number (in + out) of network packets per second for all network interfaces
total_disk_space*	The total (used + available) amount of space on all file systems in megabytes. This measurement is taken at the end of the sampling interval and includes the space held back from normal users.
total_processes*	The number of entries currently being used in the process table. This measurement is taken at the end of the sampling interval.
total_real_mem*	The total amount of real (physical) memory in megabytes. This measurement is taken at the end of the sampling interval.
total_swap_space*	The total number of megabytes available for swapping

8.11.3. TeamQuest Alert Derived Statistics

The derived statistics used by TeamQuest Alert include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQAlert.Summary
Statistic Name:	
net_errors*	The number of network errors for all network interfaces
pct_swap_free*	Percentage of unused swap space in megabytes at the end of the interval
total_processes*	The total number of processes active on the system

Section 9

Microsoft Windows Systems

Statistics for Microsoft Windows systems are collected by the TeamQuest collection agents.

This section contains a listing of the statistics collected for the system:

- System Activity Statistics (see 9.1)
- Workload Statistics (see 9.2)
- Process Statistics (see 9.3)
- Hardware Inventory Statistics (see 9.4)
- General Log Statistics (see 9.5)
- TeamQuest Log Statistics (see 9.6)
- Windows Event Log Statistics (see 9.7)
- Windows Services Statistics (see 9.8)
- Derived Statistics (see 9.9)
- Optional System Activity Statistics (see 9.10)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

9.1. System Activity Statistics

The system activity statistics are collected by the System Activity Agent and stored into the aggregation sets of the TeamQuest performance database. The statistics are classified by the hierarchy of key names. By default, to conserve CPU time and disk space, the System Activity Agent collects only a subset of all of the available metrics on your system.

The **tqbsp** version of the System Activity Agent collects the performance data using the Registry Win32 application program interface (API). The **tqw2kbsp** version of the System Activity Agent uses the Windows Management Instrumentation (WMI) API to collect the performance data. By default, the Registry API version of the System Activity Agent (**tqbsp**) is used for data collection. You can change this setting by using the Edit System Activity Agent page of the TeamQuest Manager administration interface.

While the **tqw2kbsp** and **tqbsp** parameter names employ a 5-key hierarchy (System, Class, Subclass, Statistic Name, and Resource), the Registry Win32 and WMI metrics consist of only a 4-key hierarchy (System, Object, Counter, and Instance). The **tqw2kbsp** and **tqbsp** agents create TeamQuest Manager parameter key names from Win32 and WMI API performance metric names as shown here:

TeamQuest Manager	Registry Win32 and WMI API
System	System Name
Class	Object Name
Subclass	(empty)
Statistic Name	Counter Name
Resource	Instance Name

In this section, the terms *metric* and *parameter* are used interchangeably.

Win32 Performance Data API Limitations

A number of objects are associated with the Transmission Control Protocol/Internet Protocol (TCP/IP). To enable the statistics for the TCP/IP objects, you must install the TCP/IP protocol along with the Simple Network Management Protocol (SNMP) by using the Network option in the Control Panel.

The metrics documented here represent the default sets that are collected by the System Activity Agent. Many additional metrics will be available on your system and may be collected by customizing the data collection of these agents. For more information on customizing data collection to include additional metrics, see the *TeamQuest Performance Software Administration Guide*.

Parameter Hierarchy

Class:	Cache
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/systemname/Memory
TeamQuest Table Name:	Cache
Open Table Name:	CACHE
Statistic Name:	
Copy Read Hits %	<p>The percentage of cache copy read requests that are satisfied from the cache (for example, the requests did not require a disk read to provide access to the page in the cache). A copy read is a file read operation that is satisfied by a memory copy from a cache page to the application's buffer. This method is used by the local area network (LAN) redirector for retrieving cache information, by the LAN server for small transfers, and by the disk file systems.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>

Table Field Hierarchy

Class:	CPU
Subclass:	RelativePerformance
IT Resource Name:	/TeamQuest/System/systemname/CPU
TeamQuest Table Name:	CPU.RelativePerformance
Open Table Name:	CPURELPERF
Collection interval:	1 minute
Default retentions:	1 month
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
rel_unused	The amount of CPU resources not used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
rel_used	The amount of CPU resources used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Parameter Hierarchy

Class:	CPU
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/systemname/CPU
TeamQuest Table Name:	CPU.Summary
Open Table Name:	CPUSUM
Statistic Name:	
%busy	The percentage of the elapsed time that processors are busy executing non-idle threads. This statistic can be viewed as the fraction of time spent doing useful work. This value matches the value of Hyper-V Hypervisor Logical Processor::% Total Run Time:_Total when running on Microsoft Hyper-V systems. When running on nonMicrosoft Hyper-V systems, this value matches Processor::% Processor Time:_Total. [Sequential = AVG Non-Sequential = AVG]
Class:	LogicalDisk
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/systemname/Disk
TeamQuest Table Name:	LogicalDisk
Open Table Name:	LOGDISK
Statistic Name:	
% Free Space	The ratio of the free space available on the logical disk unit to the total usable space provided by the selected logical disk drive [Sequential = AVG Non-Sequential = AVG] View Report: \report\windows\disk\pctfree.rpt
Free Megabytes	The unallocated space on the disk drive in megabytes. 1 megabyte = 1,048,576 bytes. [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\disk\dskspace.rpt
record_count	The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record. [Sequential = AVG Non-Sequential = SUM]
reduction_name	The name of reduction rule [Sequential = ID Non-Sequential = ID]
reduction_source	The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H. [Sequential = ID Non-Sequential = ID]

Microsoft Windows Systems

Class:	Memory
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Memory
Open Table Name:	MEM
Statistic Name:	
Available Bytes	<p>The amount of physical memory in bytes available to processes running on the computer. It is calculated by adding the amount of space on the Free, Zeroed, and Standby memory lists. Free memory is ready for use; Zeroed memory consists of pages of memory filled with zeros to prevent subsequent processes from seeing data used by a previous process; Standby memory is memory that has been removed from a working set of a process (its physical memory) on route to disk, but is still available to be recalled. This counter displays the last observed value only; it is not an average.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report: \report\windows\memory\freemem.rpt</p>
Cache Bytes	<p>The number of bytes currently in use by the system cache. The system cache is used to buffer data retrieved from disk or Local Area Network (LAN). The system cache uses memory not in use by active processes in the computer.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report: \report\windows\memory\cache.rpt</p>
Cache Faults/sec	<p>The number of cache faults per second. Cache faults occur whenever the cache manager does not find a file's page in the immediate cache and must ask the memory manager to locate the page elsewhere in memory or on the disk so that it can be loaded into the immediate cache.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\memory\faults.rpt</p>
Commit Limit	<p>The size in bytes of virtual memory that can be committed without having to extend the paging files. If the paging files can be extended, this is a soft limit.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Committed Bytes	<p>The size of virtual memory in bytes that has been committed as opposed to simply reserved. Committed memory must have backing storage (for example, disk) available or must be assured never to need disk storage because main memory is large enough to hold it. This is an instantaneous count, not an average over the time interval.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>

Page Faults/sec	<p>The number of page faults per second in the processor. A page fault occurs when a process refers to a virtual memory page that is not in its working set in main memory. A page fault does not cause the page to be fetched from the disk if that page is on the standby list and hence already in main memory, or if it is in use by another process that shares the page.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\memory\faults.rpt</p>
Page Reads/sec	<p>The number of times per second the disk was read to retrieve pages of virtual memory necessary to resolve page faults. Multiple pages can be read during a disk read operation.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\memory\pageops.rpt</p>
Page Writes/sec	<p>The number of times per second pages have been written to the disk because they were changed since last retrieved. Each write operation may transfer a number of pages.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\memory\pageops.rpt</p>
Pages Input/sec	<p>The number of pages read from the disk per second to resolve memory references to pages that were not in memory at the time of the reference. This counter includes paging traffic on behalf of the system cache to access file data for applications. This is an important counter to observe if you are concerned about excessive memory pressure (for example, thrashing) and the excessive paging that may result.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\memory\pagevol.rpt</p>
Pages Output/sec	<p>The number of pages that are written to disk per second because the pages have been modified in main memory</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\memory\pagevol.rpt</p>
Pages/sec	<p>The number of pages read from the disk or written to the disk per second to resolve memory references to pages that were not in memory at the time of the reference. This is the sum of Pages Input/sec and Pages Output/sec. This counter includes paging traffic on behalf of the system cache to access file data for applications. This is the primary counter to observe if you are concerned about excessive memory pressure (for example, thrashing) and the excessive paging that may result.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>

Pool Nonpaged Bytes	The number of bytes in the nonpaged pool, a system memory area where space is acquired by operating system components as they accomplish their appointed tasks. Nonpaged pool pages cannot be paged out to the paging file, but instead remain in main memory as long as they are allocated. This metric is not collected by default. [Sequential = LST Non-Sequential = SUM]
Total Bytes	The total number of bytes of physical memory [Sequential = LST Non-Sequential = SUM]
Class:	Network Interface
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	Network Interface
Open Table Name:	NETINTERFACE
Statistic Name:	
Output Queue Length	The length of the output packet queue in packets. If this is longer than 2, delays are being experienced and the bottleneck should be found and eliminated if possible. Since the requests are queued by Network Device Interface Specification (NDIS) in this implementation. The value is always 0. [Sequential = LST Non-Sequential = MAX]
Packets Outbound Errors	The number of outbound packets that could not be transmitted because of errors [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\network\errors.rpt
Packets Received Errors	The number of inbound packets that contained errors preventing the packets from being deliverable to a higher-layer protocol [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\network\errors.rpt
Packets Received Non-Unicast/sec	The number of non-unicast packets (for example, subnet broadcast or subnet multicast) delivered per second to a higher-layer protocol [Sequential = AVG Non-Sequential = SUM] View Report: \report\windows\network\non-uni.rpt
Packets Received/sec	The number of packets received per second by the network interface [Sequential = AVG Non-Sequential = SUM] View Report: \report\windows\network\packets.rpt

Packets Sent Non-Unicast/sec	<p>The number of packets requested per second to be transmitted to non-unicast addresses (for example, subnet broadcast or subnet multicast) by higher-level protocols. The number includes the packets that were discarded or not sent. [Sequential = AVG Non-Sequential = SUM] View Report: \report\windows\network\non-uni.rpt</p>
Packets Sent/sec	<p>The number of packets sent per second on the network interface [Sequential = AVG Non-Sequential = SUM] View Report: \report\windows\network\packets.rpt</p>
Class:	Objects
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/systemname/General Objects
TeamQuest Table Name:	Objects
Open Table Name:	OBJS
Statistic Name:	
Processes	<p>The number of processes in the computer at the time of data collection. This is an instantaneous count, not an average over the time interval. Each process represents the running of a program. [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\objects\active.rpt</p>
Semaphores	<p>The number of semaphores in the computer at the time of data collection. This is an instantaneous count, not an average over the time interval. Threads use semaphores to obtain exclusive access to data structures that are shared with other threads. [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\objects\active.rpt</p>
Threads	<p>The number of threads in the computer at the time of data collection. This is an instantaneous count, not an average over the time interval. A thread is the basic executable entity that can execute instructions in a processor. [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\objects\active.rpt</p>

Microsoft Windows Systems

Class:	Paging File
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /General Objects
TeamQuest Table Name:	Paging File
Open Table Name:	PAGINGFILE
Statistic Name:	
% Usage	The percentage of the page file instance in use. See also NT:Process:pagefileMB. [Sequential = AVG Non-Sequential = AVG] View Report: \report\windows\memory\pagefile.rpt
Class:	PhysicalDisk
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	PhysicalDisk
Open Table Name:	PHYSDISK
Statistic Name:	
% Disk Time	The time duration of the I/O requests. This statistic is derived as: Avg. Disk sec/Transfer * Disk Transfers/sec This statistic does not represent disk utilization. [Sequential = AVG Non-Sequential = AVG] View Report: \report\SQLServer\PctPhysicalDiskTime.rpt
% Idle Time	The percentage of time during the sampling interval that the disk was idle [Sequential = AVG Non-Sequential = AVG]
Avg. Disk Bytes/Read	The average number of bytes transferred from the disk during read operations [Sequential = AVG Non-Sequential = AVG]
Avg. Disk Bytes/Transfer	The average number of bytes transferred to or from the disk during write or read operations. This metric is not collected by default. [Sequential = AVG Non-Sequential = AVG]
Avg. Disk Bytes/Write	The average number of bytes transferred to the disk during write operations [Sequential = AVG Non-Sequential = AVG]
Avg. Disk sec/Read	The average time in seconds of a read of data from the disk [Sequential = AVG Non-Sequential = AVG] View Report: \report\windows\disk\response.rpt
Avg. Disk sec/Transfer	The average time in seconds that the disk transfer took to complete [Sequential = AVG Non-Sequential = AVG]

Avg. Disk sec/Write	<p>The average time in seconds of a write of data to the disk [Sequential = AVG Non-Sequential = AVG] View Report: \report\windows\disk\response.rpt</p>
Current Disk Queue Length	<p>The number of requests outstanding on the disk at the time the performance data is collected. It includes requests in service at the time of the snapshot. This is an instantaneous length, not an average over the time interval. Multi-spindle disk devices can have multiple requests active at one time, but other concurrent requests are awaiting service. This counter may reflect a transitory high or low queue length, but if there is a sustained load on the disk drive, it is likely that this will be consistently high. Requests are experiencing delays proportional to the length of this queue minus the number of spindles on the disks. This difference should average less than 2 for good performance. [Sequential = LST Non-Sequential = AVG]</p>
Disk Bytes/sec	<p>The number of bytes transferred per second to or from the disk during write or read operations [Sequential = AVG Non-Sequential = SUM]</p>
Disk Read Bytes/sec	<p>The number of bytes transferred per second from the disk during read operations [Sequential = AVG Non-Sequential = SUM] View Report: \report\windows\disk\volume.rpt</p>
Disk Reads/sec	<p>The number of read operations on the disk per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\windows\disk\transfer.rpt</p>
Disk Transfers/sec	<p>The number of read and write operations on the disk per second [Sequential = AVG Non-Sequential = SUM]</p>
Disk Write Bytes/sec	<p>The number of bytes transferred per second to the disk during write operations [Sequential = AVG Non-Sequential = SUM] View Report: \report\windows\disk\volume.rpt</p>
Disk Writes/sec	<p>The number of write operations on the disk per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\windows\disk\transfer.rpt</p>
record_count	<p>The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record. [Sequential = AVG Non-Sequential = SUM]</p>

reduction_name	The name of reduction rule [Sequential = ID Non-Sequential = ID]
reduction_source	The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H. [Sequential = ID Non-Sequential = ID]
Class:	Processor
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/systemname/CPU
TeamQuest Table Name:	Processor
Open Table Name:	PROC
Statistic Name:	
% DPC Time	The percentage of elapsed time that the processor spent in deferred procedure calls (DPC). When a hardware device interrupts the processor, the interrupt handler may elect to execute the majority of its work in a DPC. DPCs run at lower priority than interrupts and so permit interrupts to occur while DPCs execute. DPCs are executed in privileged mode, so this is a component of Processor:::% Privileged Time. This counter can help determine the source of excessive time being spent in privileged mode. This metric is collected by default on Microsoft Windows Server. [Sequential = AVG Non-Sequential = AVG]
% Interrupt Time	The percentage of elapsed time that the processor spent handling hardware interrupts. When a hardware device interrupts the processor, the interrupt handler executes to handle the condition, usually by signaling I/O completion and possibly issuing another pending I/O request. Some of this work may be done in a DPC (see Processor:::% DPC Time). However, time spent in DPCs is not counted as time in interrupts. Interrupts are executed in privileged mode, so this is a component of Processor:::% Privileged Time. This counter can help determine the source of excessive time being spent in privileged mode. [Sequential = AVG Non-Sequential = AVG]

% Privileged Time	<p>The percentage of processor time spent in privileged mode using non-idle threads. The Microsoft Windows service layer, the Executive routines, and the Microsoft Windows Kernel execute in privileged mode. Device drivers for most devices other than graphics adapters and printers also execute in privileged mode. Unlike some early operating systems, Microsoft Windows uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Microsoft Windows on behalf of your application may appear in other subsystem processes in addition to the Privileged Time in your process.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Reports:</p> <p>\report\windows\system\cpu-util-2000.rpt \report\windows\system\cpu-util-xp.rpt \report\windows\system\per-cpu.rpt</p>
% Processor Time	<p>A percentage of the elapsed time that a processor is busy executing a non-idle thread. It can be viewed as the fraction of the time spent doing useful work. Each processor is assigned an idle thread in the idle process that consumes those unproductive processor cycles not used by any other threads.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
% User Time	<p>The percentage of processor time spent in user mode using non-idle threads. All application code and subsystem code execute in user mode. The graphics engine, graphics device drivers, printer device drivers, and the window manager also execute in user mode. Code executing in user mode cannot damage the integrity of the Microsoft Windows Executive, Kernel, and device drivers. Unlike some early operating systems, Microsoft Windows uses process boundaries for subsystem protection in addition to the traditional protection user and privilege modes. These subsystem processes provide additional protection. Therefore, some work done by Microsoft Windows on behalf of your application may appear in other processes in addition to the Privileged Time in your process.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Reports:</p> <p>\report\windows\system\cpu-util-2000.rpt \report\windows\system\cpu-util-xp.rpt \report\windows\system\per-cpu.rpt</p>

Microsoft Windows Systems

Class:	Redirector
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /General Objects
TeamQuest Table Name:	Redirector
Open Table Name:	REDIRECTOR
Statistic Name:	
Bytes Total/sec	<p>The number of data bytes the redirector is processing per second. This includes all application data, file data, and network protocol information (such as packet headers).</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\redirect\volume.rpt</p>
Current Commands	<p>The number of requests to the redirector that are currently queued for service. If this number is much larger than the number of network adapter cards installed in the computer, the networks and the servers being accessed are seriously bottlenecked.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\redirect\queued.rpt</p>
Network Errors/sec	<p>The number of serious unexpected errors per second. These errors generally indicate the redirector and one or more servers are having serious communication difficulties. For example, a Server Manager Block (SMB) protocol error generates a network error. Network errors result in an entry in the system event log, which will provide details.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\redirect\errors.rpt</p>
Packets Received/sec	<p>The number of packets or Server Manager Blocks (SMBs) received per second by the redirector. Network transmissions are divided into packets. The average number of bytes received in a packet can be obtained by dividing Bytes Received/sec by this counter. Some packets received may not contain incoming data. For example an acknowledgment to a write made by the redirector would count as an incoming packet.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\redirect\packets.rpt</p>

Packets Transmitted/sec	<p>The number of packets or Server Manager Blocks (SMBs) sent per second by the redirector. Network transmissions are divided into packets. The average number of bytes transmitted in a packet can be obtained by dividing Redirector:::Bytes Transmitted/sec by this counter. (Bytes Transmitted/sec is not collected by default.)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\redirect\packets.rpt</p>
Server Sessions	<p>The total number of security objects the redirector is managing. For example, a logon to a server followed by a network access to the same server will establish one connection, but two sessions.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report: \report\windows\redirect\sessions.rpt</p>
Class:	Server
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/systemname/General Objects
TeamQuest Table Name:	Server
Open Table Name:	SERV
Statistic Name:	
Bytes Total/sec	<p>The number of bytes the server sent to and received from the network per second. This value provides an overall indication of the load on the server.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\server\volume.rpt</p>
Pool Nonpaged Failures	<p>The number of times allocations from nonpaged pool have failed. This indicates that the computer's physical memory is too small.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>

Microsoft Windows Systems

Class:	System
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /General Objects /TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	System
Open Table Name:	SYS
Statistic Name:	
Context Switches/sec	<p>The number of switches from one thread to another per second. Thread switches can occur either inside of a single process or across processes. A thread switch may be caused either by one thread asking another for information or by a thread being preempted by another higher priority thread becoming ready to run. Unlike some early operating systems, Microsoft Windows uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Microsoft Windows on behalf of an application may appear in other subsystem processes in addition to the Privileged Time in the application. Switching to the subsystem process causes one context switch in the application thread. Switching back causes another context switch in the subsystem thread.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\system\switches.rpt</p>
Processor Queue Length	<p>The instantaneous length of the processor queue in units of threads. This counter is always 0 unless you are also monitoring a thread counter. All processors use a single queue in which threads wait for processor cycles. This length does not include the threads that are currently executing. A sustained processor queue length greater than 2 generally indicates processor congestion. This is an instantaneous count, not an average over the time interval.</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report: \report\windows\system\cpu-qlen.rpt</p>
System Calls/sec	<p>The frequency of calls to Microsoft Windows system service routines per second. These routines perform all of the basic scheduling and synchronization of activities on the computer and provide access to non-graphical devices, memory management, and name space management.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\system\syscalls.rpt</p>

Class:	TCP
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	TCP
Open Table Name:	TCP
Statistic Name:	
Connections Active	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\tcp\connects.rpt
Connections Established	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\tcp\connects.rpt
Connections Passive	The number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state [Sequential = LST Non-Sequential = SUM] View Report: \report\windows\tcp\connects.rpt
Class:	TCPv4
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	TCPv4
Open Table Name:	TCPV4
Statistic Name:	
Connections Active	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state [Sequential = LST Non-Sequential = SUM]
Connections Established	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT [Sequential = LST Non-Sequential = SUM]
Connections Passive	The number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state [Sequential = LST Non-Sequential = SUM]

Microsoft Windows Systems

Class:	TCPv6
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	TCPv6
Open Table Name:	TCPV6
Statistic Name:	
Connections Active	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state [Sequential = LST Non-Sequential = SUM]
Connections Established	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT [Sequential = LST Non-Sequential = SUM]
Connections Passive	The number of times TCP connections have made a direct transition to the SYN-RCVD state from the LISTEN state [Sequential = LST Non-Sequential = SUM]

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
bsp interval	The number of seconds elapsed between two data samples of the System Activity Agent [Sequential = SUM Non-Sequential = ID]
tqbsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqbsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

9.2. Workload Statistics

TeamQuest Manager maintains workload statistics in the performance database. All statistics listed here are collected by the Process-Workload Agent into workloads in the aggregation sets defined for the database.

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
twarp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
twarp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]
warp interval	The number of seconds elapsed between two data samples of the Process-Workload Agent [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

Parameter Hierarchy

Class:	Workload
Subclass:	by Workload
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Workload/ <i>workloadset</i> / <i>workload</i>
TeamQuest Table Name:	Workload.by Workload
Open Table Name:	WLBYWORKLOAD
Workload Set:	WLS1, WLS2, ...
Workload:	WL1, WL2, ...
Statistic Name:	
%cpu	<p>The percentage of total CPU consumed by the workload. Total CPU time is the sampling interval times the number of CPUs on the system. Thus, if the sum of %cpu for all workloads is less than 100%, then that missing usage is not accounted for in any workload.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\windows\workload\cpu-util.rpt</p>
databytes	<p>The number of bytes read or written by all of the processes in the workload in I/O operations during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
dataops	<p>The number of read and write I/O operations issued by all of the processes in the workload during the sampling interval. I/O activity generated by the processes includes file, network and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
elapsed	<p>The sum of the elapsed time in seconds of all of the processes in the workload. Dividing this number by the number of processes in the workload (nproc) gives the average time a process in the workload existed during the sampling interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
handles	<p>The total number of handles currently open by all of the processes in the workload</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
otherbytes	<p>The number of bytes issued by all of the processes in the workload to I/O operations during the sampling interval that are neither a read nor a write operation. An example of this type of operation would be a control function. I/O activity generated by the processes includes file, network and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
otherops	<p>The number of I/O operations issued by all of the processes in a workload during the sampling interval that are neither a read nor a write operation. An example of this type of operation would be a control function. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>

pagefaults	<p>The sum of the page faults that occurred in all of the threads of the workload</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\workload\pageflts.rpt</p>
pagefileMB	<p>The current number of megabytes (MB) the process has used in the paging files. Paging files store pages of memory used by the process that are not contained in other files. Paging files are shared by all of the processes. Lack of space in paging files can prevent other processes from allocating memory.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
pcompleted	<p>The number of processes that end in the sampling interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
pongoing	<p>The number of ongoing (active) processes at the end of the sampling interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\workload\ndproc.rpt</p>
privateMB	<p>The current number of megabytes (MB) the process has allocated that cannot be shared with other processes</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
privcpu	<p>The privileged mode CPU time in seconds used by the workload. Privileged mode CPU time is the time spent in kernel mode (for example, the time spent in executing system calls, paging, and so on).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\workload\privcpu.rpt</p>
pstarted	<p>The number of processes that started in the sampling interval and were still active at the end of the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
readbytes	<p>The number of reads from I/O operations by all of the processes in the workload during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
readops	<p>The number of read operations issued by the processes in the workload during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
record_count	<p>The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
reduction_name	<p>The name of reduction rule</p> <p>[Sequential = ID Non-Sequential = ID]</p>

reduction_source	<p>The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
tcompleted	<p>The number of threads that completed for all of the processes in the workload in the sampling interval. The value represents a sum of the differences in active threads from the previous interval for each process in the workload. If fewer threads are active for an individual process at the end of the current sampling interval than in the previous interval, the value will be nonzero for that process. Otherwise, the value will be zero.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
tongoing	<p>The number of threads active for all of the processes in the workload at the end of the sampling</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\workload\ntproc.rpt</p>
totcpu	<p>The total CPU time in seconds used by all of the processes in the workload in the current interval. The value is the sum of usercpu and privcpu.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\workload\totcpu.rpt</p>
tstarted	<p>The number of threads that all of the processes in the workload started in the sampling interval. The value represents a sum of the differences in active threads from the previous interval for each process in the workload. If more threads are active for an individual process at the end of the current sampling interval than in the previous interval, the value will be nonzero for that process. Otherwise, this value will be zero.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
usercpu	<p>The user CPU time in seconds used by the workload. User CPU time is the time CPU spent running in user mode.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\windows\workload\usercpu.rpt</p>
virtualMB	<p>The virtual address space size in megabytes (MB) of all of the running processes in the workload at the end of the sampling interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
writebytes	<p>The number of bytes written to I/O operations by all of the processes in the workload during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>

writeops	The number of write I/O operations issued by all of the processes in the workload during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os. [Sequential = SUM Non-Sequential = SUM]
wssMB	The working set size in megabytes (MB) occupied by all of the running processes in the workload at the end of the sampling interval [Sequential = AVG Non-Sequential = SUM] View Report: \\report\\windows\\workload\\wss.rpt

9.3. Process Statistics

The Process-Workload Agent (**tqwarp** or **tqw2kwarp**) collects process data, calculates the usage of every process in a given interval, applies the reduction definitions to each process, and stores the reduced process data. It also applies workload definitions to the reduced process data and stores system resource usage by workload.

Reduced Process Records

The Process-Workload Agent collects process data and reduces the data according to the user-defined reduction definitions. A reduction definition may cause multiple processes to be merged into a single process record. Thus, a process record contains data about one or more processes. When you are looking at the resource usage numbers, it is important to know how many processes a process record actually represents. The nproc data item indicates exactly how many processes each process record is representing. When a process record is representing more than one process, the resource usage fields such as wssMB, usercpu, and pagefaults are the sum of the resource usage of the individual processes. When all of the processes do not have the same value for a field, the identifier fields such as command are set to <Multi>. When data for some fields is not available, the fields are set to <N/A>.

Disabling Reduction Definitions

If you want to look at the details of every individual process and do not wish to have merged process records, you must disable reduction processing by making all reduction sets inactive. However, with reduction processing disabled, more records have to be stored and more disk space is needed. For information on disabling reductions, see the *TeamQuest Performance Software Administration Guide*.

Retrieving Hardware Configuration Information

The Process-Workload Agent retrieves hardware configuration information. The information is stored upon startup and once-a-day in the HINV.Summary and the HINV.Devices table files of the TeamQuest performance database. The information is also stored if the agent detects a change in configuration.

Table Field Hierarchy

Class:	NT
Subclass:	Process
IT Resource Name:	/TeamQuest/System/systemname /TeamQuest/System/systemname/Process
TeamQuest Table Name:	NT.Process
Open Table Name:	NTPROC
Collection interval:	Based on the primary aggregation set
Default retention:	1 day
Table type:	Performance

Note: *The collection interval depends on the Processes Only setting in the configuration file for the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.*

Statistic Name	Description
actual_interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
btime	The start time of the process. For process records representing more than one process, the field shows the earliest of the start times. [Sequential = FST Non-Sequential = ID]
command	The command name of the process. Up to 51 characters are displayed. [Sequential = ID Non-Sequential = ID]
cproc	The number of processes that ended in the sampling interval [Sequential = SUM Non-Sequential = SUM]
cthread	The number of threads that completed in the process in the sampling interval. The value represents the difference in active threads from the previous interval. If fewer threads are active at the end of the current sampling interval than in the previous interval, the value will be nonzero. Otherwise, this value will be zero. [Sequential = SUM Non-Sequential = SUM]
databytes	The number of bytes read or written by the process in I/O operations during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os. [Sequential = SUM Non-Sequential = SUM]

databytes_t	The number of bytes read or written by the process in I/O operations since the process started running. I/O activity generated by the processes includes file, network, and device I/Os. This statistic is not collected by default. [Sequential = LST Non-Sequential = SUM]
dataops	The number of read and write I/O operations issued by the process during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os. [Sequential = SUM Non-Sequential = SUM]
dataops_t	The number of read and write I/O operations issued since the process started running. I/O activity generated by the processes includes file, network, and device I/Os. This statistic is not collected by default. [Sequential = LST Non-Sequential = SUM]
elapsed	The elapsed time in seconds for the current interval. This number tells how long a process existed in the current interval. [Sequential = SUM Non-Sequential = SUM]
elapsed_t	The total elapsed time in seconds. This number tells how long a process existed since it started. For a single process, this is the same as timestamp minus btime. [Sequential = LST Non-Sequential = SUM]
fullcmd	The full command line of the process. Up to 2048 characters are displayed. This statistic is not collected by default. When this statistic is disabled, <NA> is displayed. [Sequential = ID Non-Sequential = ID]
handles	The total number of handles currently open by the process. This number is the sum of the handles currently open by each thread in the process. [Sequential = LST Non-Sequential = SUM]
Interval	The desired interval for the data collection which may differ from the actual_interval value [Sequential = SUM Non-Sequential = ID]
job_name	The job name associated with the process. If no job name is associated with the process or if a job name collection has been disabled, this field will be blank.
nonp_poolMB	The number of megabytes (MB) in the nonpaged pool at the end of the sampling interval. The nonpaged pool is a system memory area where space is acquired by the operating system components as they accomplish their appointed tasks. Nonpaged pool pages cannot be paged out to the paging file, but instead remain in main memory as long as they are allocated. This statistic is not collected by default. [Sequential = AVG Non-Sequential = SUM]
nproc	The number of processes the process record represents. In a reduced record, it is the number of processes that were merged together to form a single process record. [Sequential = AVG Non-Sequential = SUM]
oproc	The number of ongoing processes at the end of the sampling interval [Sequential = LST Non-Sequential = SUM]

otherbytes	<p>The number of bytes issued by the process to I/O operations during the sampling interval that are neither a read nor a write operation. An example of this type of operation would be a control function. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
otherbytes_t	<p>The number of bytes issued by the process to I/O operations since the process started running that are neither a read nor a write operation. An example of this type of operation would be a control function. I/O activity generated by the processes includes file, network, and device I/Os. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
otherops	<p>The number of I/O operations issued by the process during the sampling interval that are neither a read nor a write operation. An example of this type of operation would be a control function. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
otherops_t	<p>The number of I/O operations issued by the process since the process started running that are neither a read nor a write operation. An example of this type of operation would be a control function. I/O activity generated by the processes includes file, network, and device I/Os. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
othread	<p>The number of threads active in the process at the end of the interval. An instruction is the basic unit of execution in a processor and a thread is the object that instructs. Every running process has at least one thread.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
p_poolMB	<p>The number of megabytes (MB) in the paged pool at the end of the sampling interval. The paged pool is a system memory area where space is acquired by the operating system components as they accomplish their appointed tasks. Paged pool pages can be paged out to the paging file when not accessed by the system for sustained periods of time. This statistic is not collected by default.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
pagefaults	<p>The number of page faults for the process during the sampling interval. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory. This will not cause the page to be fetched from disk if it is on the standby list and is already in main memory, or if it is in use by another process that shares the page.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
pagefaults_t	<p>The number of page faults since the process started running. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>

pagefileMB	<p>The current number of megabytes (MB) the process has used in the paging files. Paging files store pages of memory used by the process that are not contained in other files. Paging files are shared by all of the processes. Lack of space in paging files can prevent other processes from allocating memory.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
pagefileMBpeak	<p>The maximum number of megabytes (MB) the process has used in the paging files. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
pctcpu	<p>The percentage of elapsed time that all of the threads of the process used the processor to execute instructions. An instruction is the basic unit of execution in a computer. A thread is the object that executes instructions. A process is the object created when a program is run. Code executed to handle certain hardware interrupts or trap conditions may be counted for the process.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
pid	<p>The unique identifier of the process. The pid numbers are reused, so the number only identifies a process for the lifetime of that process.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
priority	<p>The current base priority of the process. Threads within a process can raise and lower their base priority relative to the base priority of the process. If the process record represents more than one process, the priority is an average of the priority values of the constituent processes.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
privateMB	<p>The current number of megabytes (MB) the process has allocated that cannot be shared with other processes. This statistic is not collected by default.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
privcpu	<p>The number of seconds in the interval that the process threads have spent executing code in privileged mode. When a Microsoft Windows system service is called, the service often runs in privileged mode to gain access to system-private data. Such data is protected from access by threads executing in user mode. Calls to the system may be explicit or they may be implicit, such as when a page fault or an interrupt occurs. Unlike some early operating systems, Microsoft Windows uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Microsoft Windows on behalf of your application may appear in other subsystem processes in addition to the Privileged Time in your process.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
privcpu_t	<p>The number of seconds that the process threads have spent executing code in privileged mode since the process started running. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>

readbytes	<p>The number of bytes read from I/O operations by the process during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
readbytes_t	<p>The number of bytes read from I/O operations by the process since the process started running. I/O activity generated by the processes includes file, network, and device I/Os. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
readops	<p>The number of read I/O operations issued by the process during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
readops_t	<p>The number of read I/O operations issued by the process since the process started running. I/O activity generated by the processes includes file, network, and device I/Os. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
redname	<p>The reduction name of the process record. If a process does not match any of the reduction definitions, then it is not reduced and the field will be blank.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
sample_end_time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
sproc	<p>The number of processes that started in the sampling interval and were still active at the end of the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
sthread	<p>The number of threads that the process started in the sampling interval. The value represents the difference in active threads from the previous interval. If more threads are active at the end of the current sampling interval than in the previous interval, the value will be nonzero. If the process started in the current sampling interval, the value will be equal to the number of active threads at the end of the interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
totcpu	<p>The total CPU time in seconds used by the process record in the current interval. This number is the same as the sum of usercpu and privcpu.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>

totcpu_t	<p>The total CPU time in seconds used by the process record since it started. This number is the same as the sum of usercpu_t and privcpu_t.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
user	<p>The name of the user running the process. This information may not be available for all of the processes.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
usercpu	<p>The number of seconds in the interval that the process threads have spent executing in user mode. Applications execute in user mode, as do subsystems, like the window manager and the graphics engine. Code executing in user mode cannot damage the integrity of the Microsoft Windows Executive, Kernel, and device drivers. Unlike some early operating systems, Microsoft Windows uses process boundaries for subsystem protection in addition to the traditional protection of user and privileged modes. These subsystem processes provide additional protection. Therefore, some work done by Microsoft Windows on behalf of your application may appear in other subsystem processes in addition to the Privileged Time in your process.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
usercpu_t	<p>The number of seconds that the process threads have spent executing code in user mode since the process started running. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
virtualMB	<p>The current size in megabytes (MB) of the virtual address space that the process is using. The use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is, however, finite and by using too much, the process may limit its ability to load libraries.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
virtualMBpeak	<p>The maximum number of megabytes (MB) of virtual address space the process has used at any one time. This statistic is not collected by default.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>

Workload	<p>The workload set and the workload associated with the process. When the Process-Workload Agent stores the process record, this field is blank. When process records are reported, the workload can be evaluated and is shown in the report.</p> <p>This field is available for reporting only when using TeamQuest Analyzer and TeamQuest tView.</p> <p>Workload evaluation takes place when data is retrieved from the database, based on workload sets defined in the database where the data is stored. Workload sets reported in the Workload field do not have to be active.</p> <p>For more information on workload evaluation, see the <i>TeamQuest Analyzer User Guide</i> or the <i>TeamQuest Performance Software Command Line Interfaces Reference Manual</i>. [Sequential = ID Non-Sequential = ID]</p>
workload:wlsname	<p>There is one field for each <i>wlsname</i> (Workload Set Name). The value for this field shows the name of the workload to which the process belongs. If a process belongs to none of the workloads, it will display the workload name "OTHER."</p> <p>This field is available for reporting only when using TeamQuest View or TeamQuest cView. [Sequential = ID Non-Sequential = ID]</p>
writebytes	<p>The number of bytes written to I/O operations by the process during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os. [Sequential = SUM Non-Sequential = SUM]</p>
writebytes_t	<p>The number of bytes written to I/O operations by the process since the process started running. I/O activity generated by the processes includes file, network, and device I/Os. This statistic is not collected by default. [Sequential = LST Non-Sequential = SUM]</p>
writeops	<p>The number of write I/O operations issued by the process during the sampling interval. I/O activity generated by the processes includes file, network, and device I/Os. [Sequential = SUM Non-Sequential = SUM]</p>
writeops_t	<p>The number of write I/O operations issued by the process since the process started running. I/O activity generated by the processes includes file, network, and device I/Os. This statistic is not collected by default. [Sequential = LST Non-Sequential = SUM]</p>

wssMB	The size in megabytes (MB) of the working sets at the end of the interval. The working set is the set of memory pages recently touched by the thread in the process. If free memory in the computer is above a threshold, pages are left in the working set of a process even if they are not in use. When free memory falls below the threshold, pages are trimmed from working sets. If the pages are needed, they will be soft-faulted back into the working set before they leave main memory. [Sequential = AVG Non-Sequential = SUM]
wssMBpeak	The maximum number of megabytes (MB) in the working set of the process at any point in time. This statistic is not collected by default. [Sequential = LST Non-Sequential = SUM]

9.4. Hardware Inventory Statistics

The Process-Workload Agent (**tqwarp** or **tqw2kwarp**) retrieves hardware configuration information. The information is stored upon startup and once-a-day in the HINV.Summary, HINV.CPUModel, HINV.CPU Thread Speeds, HINV.Devices, HINV.FileSyste, and the HINV.FileSystemToDevice table files of the TeamQuest performance database. The information is also stored if the agent detects a change in configuration.

Note: *The storage of hardware inventory records depends on the Hardware Inventory setting in the configuration file of the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.*

Table Field Hierarchy

Class:	HINV
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	HINV.Summary
Open Table Name:	HINVSUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
core_multi_thread	The status or ability of the processor to support multiple independent threads. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cores_per_chip	The number of cores or processors on an individual chip. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]

Microsoft Windows Systems

cpu_chips	The number of CPU chips or sockets. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cpu_count	The number of configured processors. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cpu_speed	The speed of the processor in MHz or GHz [Non-Sequential = ID]
cpu_type	The basic instruction set architecture of the current system [Non-Sequential = ID]
logical_cpu_count	The number of logical processors [Non-Sequential = ID]
mem_size	The size of configured random access memory in kilobytes, where 1 kilobyte = 1,024 bytes [Non-Sequential = ID]
memory	The size of configured random access memory in megabytes, where 1 megabyte = 1,048,576 bytes [Non-Sequential = ID]
memory_size	The size of configured random access memory in megabytes or gigabytes [Non-Sequential = ID]
model	Name of the hardware implementation or platform [Non-Sequential = ID]
os_release	The name and level of this implementation of the operating system [Non-Sequential = ID]
pagesize	The size of a page of memory [Non-Sequential = ID]
partition_type	The partition type of the system. The value indicates the system hypervisor type, guest type, logical partition type, zone type, or logical domain type. If the system does not have a partition type, this field will be blank. [Non-Sequential = ID]
serial	The hardware-specific serial number of the physical machine [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_identifier	Information used to identify the system [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]

timezone The time zone where the data was collected
[Non-Sequential = ID]

TQLevel The level of TeamQuest Manager
[Non-Sequential = ID]

The HINV.CPUModel table stores best-match, relative performance data about the system configuration. This table is created by the Hardware Inventory Agent (**tqhinv**) to map physical hardware to a CPU model that describes performance in relative terms. This table is not created for any virtualized system. It is populated for physical systems only. It is not populated for VMware guests, Hyper-V guests, Solaris logical domains (LDOMs), Solaris guest LDOMs, KVM guests, and Linux on POWER systems.

Table Field Hierarchy

Class: HINV
Subclass: CPUModel
IT Resource Name: /TeamQuest/System/systemname
TeamQuest Table Name: HINV.CPUModel
Open Table Name: HINVCPU
Collection interval: N/A
Default retention: 1 year
Table type: State

Statistic Name	Description
cpu_chips	The number of CPU chips or sockets [Non-Sequential = ID]
cpu_confidence	The percentage of confidence in the correctness of the CPU match based on model, frequency, and configuration (chips, cores, threads) [Non-Sequential = SUM]
cpu_cores	The number of CPU cores or processors on an individual CPU chip [Non-Sequential = ID]
cpu_name	The name of the selected CPU [Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Non-Sequential = ID]
cpu_speed	The speed of the processor in megahertz (MHz) or gigahertz (GHz) [Non-Sequential = ID]
cpu_threads	The number of CPU threads on an individual CPU core or processor [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]

Time	The timestamp of the data sample [Non-Sequential = LST]
user_override	The user override status of the default TeamQuest generated CPU match. This field is not currently used and should appear as 0. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	CPU Thread Speeds
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.CPU Thread Speeds
Open Table Name:	HINVCPU THREADSPEEDS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
speed_up_factor	The performance improvement when there are multiple active threads per core, compared to when there is only one active thread per core [Sequential = AVG Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
thread_number	The number of active threads [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = ID Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	Devices
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.Devices
Open Table Name:	HINVDEVS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
class	The device classification: controller, disk, or tape [Non-Sequential = ID]
controller	The device path indicator which defines a connection to another device [Non-Sequential = ID]
lun_id	The globally unique Logical Unit Number (LUN) identifier for Storage Area Network (SAN) based disk devices. This field is blank for non-SAN based disk devices, CD-ROM drives, tape drives, and so on. [Non-Sequential = ID]
name	The unique identifier for this device [Non-Sequential = ID]
name2	The alternate device name. The field may be blank. [Non-Sequential = ID]
product	The product identifier. This field may be blank. [Non-Sequential = ID]
revision	The revision level for this product. This field may be blank. [Non-Sequential = ID]
rpm	The speed at which the media spins. If an actual value cannot be obtained for the device, a default value of 7,200 is used. [Non-Sequential = ID]
sequence	The sequence number of the device [Non-Sequential = ID]
swap	A true or false statement indicating whether or not there is a swap file on the device [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]

Microsoft Windows Systems

Time	The timestamp of the data sample [Non-Sequential = ID]
vendor	The name of the device vendor. This field may be blank. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	FileSystem
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.FileSystem
Open Table Name:	HINVFILESYS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
BlkSize	The size of a block on the file system [Non-Sequential = ID]
Device	The logical disk containing the file system [Non-Sequential = ID]
Name	The unique identifier for the file system [Non-Sequential = ID]
Source	The source physical disk or logical volume of the file system. This field is always blank for this platform. [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
TotBlks	The total number of blocks on the file system [Non-Sequential = ID]
TotFiles	The maximum total number of files, as represented by inodes, possible on the file system. Some inodes may be used for entities other than visible files. [Non-Sequential = ID]
TotSize	The total amount of space on the file system in megabytes [Non-Sequential = ID]
Type	The type of the file system [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	FileSystemToDevice
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.FileSystemToDevice
Open Table Name:	HINVFSTODEV
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
Device	The name of the device [Non-Sequential = ID]
FileSystem	The name of the file system [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

9.5. General Log Statistics

The General Log Agent (**tqglp**) is used to collect log messages generated by application programs. The General Log Agent stores the messages in the TeamQuest performance database for analysis. Examples include backup, security, database, and Web server applications.

Table Field Hierarchy

Class:	System
Subclass:	General Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /General Log
TeamQuest Table Name:	System.General Log
Open Table Name:	SYSGENERALLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Message	The message text [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The message type [Non-Sequential = ID]

9.6. TeamQuest Log Statistics

The following statistics are stored in the performance database tables by the TeamQuest Log Agent (**tqlog**). The collection interval and retention periods can be modified. For more information on modifying the collection interval and retention periods, see the *TeamQuest Performance Software Administration Guide*.

Table Field Hierarchy

Class:	Service
Subclass:	TeamQuest Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /TeamQuest Log
TeamQuest Table Name:	Service.TeamQuest Log
Open Table Name:	SVCTQLOG
Collection interval:	N/A
Default retention:	1 day
Table type:	Event

Statistic Name	Description
Filename	The name of the TeamQuest log file that was the source of the message text [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The log message type. This is always set to tqlog . [Non-Sequential = ID]

9.7. Windows Event Log Statistics

The Windows Event Log Agent (**tqw2kevent**) collects information on Microsoft Windows event log records.

Table Field Hierarchy

Class:	NT
Subclass:	Event
IT Resource Name:	/TeamQuest/System/systemname/WindowsEventLog
TeamQuest Table Name:	NT.Event
Open Table Name:	NTEVENT
Collection interval:	N/A
Default retention:	6 months
Table type:	Event

Statistic Name	Description
Category	A numeric categorization of the event. This is supplied by the source of the event. [Non-Sequential = ID]
Count	The number of events the record represents. This count is useful when event records are consolidated, but is usually one. [Non-Sequential = ID]
Data	The string representation of the binary data that accompanied the event. This field may be blank. [Non-Sequential = ID]
Event	The event identifier value as displayed by the Windows event viewer application [Non-Sequential = ID]
Event_Computer	The name of the computer that generated the event [Non-Sequential = ID]
Event_Time	The time when the event source generated the event [Non-Sequential = ID]
Event_Timestamp	The timestamp of when the event source generated the event [Non-Sequential = ID]
Log_File	The name of the Windows event log file [Non-Sequential = ID]
Message	The event message as it appears in the Windows event log [Non-Sequential = ID]
Sequence	A numeric sequence number that, along with the System and Time fields, uniquely identifies the event record [Non-Sequential = ID]
Source	The name of the source (application, service, driver, or subsystem) that generated the event [Non-Sequential = ID]

System	The name of the system where the agent collecting the event data is running. The System field, along with the Sequence and Time fields, uniquely identifies the event record. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The type of the event [Non-Sequential = ID]
User	The user name of the logged-on user when the event occurred, if possible to determine [Non-Sequential = ID]

9.8. Windows Services Statistics

The Windows Services Agent (**tqw2kserv**) collects information on the services registered with the Windows service control manager.

Table Field Hierarchy

Class:	NT
Subclass:	Service
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Windows Services
TeamQuest Table Name:	NT.Service
Open Table Name:	NTSVC
Collection interval:	Based on the default aggregation set
Default retention:	3 days
Table type:	Performance

Statistic Name	Description
Caption	A one-line description of the service. It is also known as the external name of the service which is visible from the Windows Services control panel application. [Sequential = LST Non-Sequential = ID]
Count	The number of services the record represents. This count is useful when event records are consolidated, but is usually one. [Sequential = SUM Non-Sequential = SUM]
ExitCode	The Win32 error code defining any problems encountered in starting or stopping the service [Sequential = LST Non-Sequential = ID]
Interval	The expected collection interval that the agent is sampling [Sequential = SUM Non-Sequential = ID]

Name	<p>The internal name of the service that uniquely identifies the service and provides an indication of the functionality that is managed. Name, along with the Time and System fields uniquely identifies the record.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
ProcessID	<p>The process identifier of the service</p> <p>[Sequential = LST Non-Sequential = ID]</p>
StartMode	<p>The start mode for the service</p> <p>[Sequential = LST Non-Sequential = ID]</p>
State	<p>The current state of the service</p> <p>[Sequential = LST Non-Sequential = ID]</p>
System	<p>The name of the system where the agent collecting the data is running. The System field, along with the Name and Time fields, uniquely identifies the record. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>

9.9. Derived Statistics

Some products within TeamQuest Performance Software use derived statistics to display common statistics across different platforms. The derived statistics are inserted into the performance database when databases are created. In this section, a derived statistic is marked with an asterisk (*).

You can find information on the following derived statistics:

- TeamQuest Alert Derived Statistics
- TeamQuest On the Web Derived Statistics
- TeamQuest View Derived Statistics

9.9.1. TeamQuest Alert Derived Statistics

The derived statistics used by TeamQuest Alert include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQAlert.Summary
Statistic Name:	
free_physical_mem*	The amount of physical memory available to processes in megabytes
free_virtual_mem*	The size of the virtual memory in megabytes currently on the zeroed, free, and standby lists. Zeroed and free memory is ready for use, with zeroed memory cleared to zeros. Standby memory is memory removed from a process' working set but is still available. This is an instantaneous count, not an average over the time interval.
paging_file_pct_free*	The percentage of all the page file instances that are free. See also NT.Process.pagefileMB.
pkt_errors*	The number of received and outbound packets for all network interfaces, which contain errors preventing the packets from being delivered
pkts_per_sec*	The number of packets are sent and received for all network interfaces
total_processes*	The total number of processes active on the system

9.9.2. TeamQuest On the Web Derived Statistics

The derived statistics used by TeamQuest On the Web include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQWeb.Summary
Statistic Name:	
disk_xfers_per_sec*	The number of read and write operations for all of the disks
free_disk_space*	The unallocated space for all of the disk drives in megabytes, where 1 megabyte = 1,048,576 bytes
free_physical_mem*	The amount of physical memory available to processes in megabytes
free_virtual_mem*	The size of the virtual memory in megabytes currently on the zeroed, free, and standby lists. Zeroed and free memory is ready for use, with zeroed memory cleared to zeros. Standby memory is memory removed from a process' working set but is still available. This is an instantaneous count, not an average over the time interval.
paging_file_pct_free*	The percentage of all of the page file instances that are free
pct_disk_busy*	The percentage of elapsed time that all of the disk drives were busy servicing read or write requests
pct_priv_cpu*	The percentage of non-idle processor time spent in privileged mode. Privileged mode is a processing mode designed for operating system components and hardware-manipulating drivers. It allows direct access to hardware and all of the memory. The alternative, user mode, is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. The operating system switches application threads to privileged mode to access operating system services. The percentage of privileged time includes time servicing interrupts and Deferred Procedure Calls (DPC). A high rate of privileged time might be attributable to a large number of interrupts generated by a failing device. This counter displays the average busy time as a percentage of the sample time.
pct_proc_cpu*	The percentage of time that the processor is executing a non-idle thread. This counter was designed as a primary indicator of processor activity. It is calculated by measuring the time that the processor spends executing the thread of the idle process in each sampling interval, and subtracting that value by 100% (each processor has an idle thread which consumes cycles when no other threads are ready to run). It can be viewed as the percentage of the sampling interval spent doing useful work. This counter displays the average percentage of busy time observed during the sampling interval.

pct_usr_cpu*	The percentage of non-idle processor time spent in user mode. User mode is a restricted processing mode designed for applications, environment subsystems, and integral subsystems. The alternative, privileged mode, is a processing mode designed for operating system components and hardware-manipulating drivers. It allows direct access to hardware and all of the memory. The operating system switches application threads to privileged mode to access operating system services. This counter displays the average busy time as a percentage of the sample time.
pkt_errors*	The number of received and outbound packets for all of the network interfaces that contain errors preventing the packets from being delivered
pkts_in_per_sec*	The number of packets received per second by all of the network interfaces
pkts_out_per_sec*	The number of packets sent per second by all of the network interfaces
pkts_per_sec*	The number of packets sent and received per second for all of the network interfaces
total_disk_space*	The allocated space for all of the disk drives in megabytes, where 1 megabyte = 1,048,576 bytes.

9.9.3. TeamQuest View Derived Statistics

The derived statistics used by TeamQuest View include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQView.Summary
Statistic Name:	
Free Physical Memory (MB)*	The amount of physical memory available to processes in megabytes (MB) View Report: \report\windows\memory\physmem.rpt
Free Virtual Memory (MB)*	The amount of virtual memory available to processes in megabytes (MB) View Report: \report\windows\memory\virtmem.rpt
Total Physical Memory (MB)*	The amount of physical, random access in megabytes (MB) View Report: \report\windows\memory\physmem.rpt
Total Virtual Memory (MB)*	The amount of virtual memory that can be committed to all of the processes without enlarging the paging file in megabytes (MB) View Report: \report\windows\memory\virtmem.rpt

9.10. Optional System Activity Statistics

The system activity statistics are collected by the System Activity Agent (**tqbsp** or **tqw2kbsp**). By default, to conserve CPU time and disk space, the System Activity Agents collect only a subset of the available metrics on your system.

Some optional statistics that can be collected from Microsoft Exchange 2000, Microsoft Exchange 5.5, Microsoft Internet Information Server (IIS), and Microsoft SQL Server. These statistics are collected by the System Activity Agent directly from the Microsoft Windows registry.

To have the System Activity Agent collect these optional statistics, you will need to turn on the collection of the appropriate statistic in the configuration of the System Activity Agent. You do this through the TeamQuest Manager browser interface. For information on how to turn on data collection for individual statistics, see the *TeamQuest Performance Software Administration Guide*.

Note: *Some of the optional statistics available are user statistics. User statistics can only be collected if the TeamQuest Manager service is configured under that user. For example, if the TeamQuest Manager service is configured under the local system account, Microsoft Outlook statistics are not collected because Microsoft Outlook statistics require user authentication. However, if TeamQuest Manager is configured with a user account, Microsoft Outlook statistics are collected but other statistics that require a local system account are not collected.*

The statistics listed in the following subsections are the statistics reported in the default reports that are distributed with TeamQuest View. For more information on other available statistics, see the respective third-party documentation.

You can find information on the following optional system activity statistics:

- Microsoft Exchange 2000
- Microsoft Exchange 5.5
- Microsoft IIS Statistics
- Microsoft SQL Server

9.10.1. Microsoft Exchange 2000

The following statistics can be collected by the System Activity Agent from Microsoft Exchange 2000.

Parameter Hierarchy

Class:	MSExchangeIS
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /MSExchange
TeamQuest Table Name:	MSExchangeIS
Open Table Name:	MSEXIS
Statistic Name:	
Active Connection Count	<p>The number of connections that have shown some activity in the last 10 minutes [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release2000\ISActiveConnectionAndUserCount.rpt</p>
Active User Count	<p>The number of user connections that have shown some activity in the last 10 minutes [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release2000\ISActiveConnectionAndUserCount.rpt</p>
Virus Scan Files Cleaned	<p>The number of files that have been scanned for viruses and were cleaned [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release2000\VirusScannedFilesCleanedAndQuarantined.rpt</p>
Virus Scan Files Quarantined	<p>The number of files that have been scanned for viruses and were quarantined [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release2000\VirusScannedFilesCleanedAndQuarantined.rpt</p>
Virus Scan Messages Cleaned/sec	<p>The number of messages that have been scanned for viruses and were cleaned per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release2000\VirusScannedMsgsCleanedAndQuarantined.rpt</p>

Microsoft Windows Systems

Virus Scan Messages Quarantined/sec	<p>The number of messages that have been scanned for viruses and were quarantined per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftExchg\Release2000\ VirusScannedMsgsCleanedAndQuarantined.rpt</p>
Class:	MSExchangeIS Mailbox
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /MSExchange
TeamQuest Table Name:	MSExchangeIS Mailbox
Open Table Name:	MSEXISMAILBOX
Resource:	mailbox type1, mailbox type2, ...
Statistic Name:	
Average Local Delivery Time	<p>The average time between the submission of a message and the delivery to all local recipients (recipients on the same server) for the last 10 messages</p> <p>[Sequential = LST Non-Sequential = AVG]</p> <p>View Report:</p> <p>\report\MicrosoftExchg\Release2000\ ISMailboxAvgLocalDeliveryTime.rpt</p>
Message Recipients Delivered/min	<p>The number of messages that are delivered to recipients per minute</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftExchg\Release2000\ ISMailboxMsgsAndRecipientsDeliveredperMin.rpt</p>
Messages Delivered/min	<p>The number of messages that are delivered to all recipients per minute</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftExchg\Release2000\ ISMailboxMsgsAndRecipientsDeliveredperMin.rpt</p>
Receive Queue Size	<p>The number of messages in the private information stores receive queue</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftExchg\Release2000\ ISMailboxSendAndReceiveQueueSize.rpt</p>
Send Queue Size	<p>The number of messages in the information stores send queue</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftExchg\Release2000\ ISMailboxSendAndReceiveQueueSize.rpt</p>

Class:	MSExchangeMTA Connections
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /MSExchange
TeamQuest Table Name:	MSExchangeMTA Connections
Open Table Name:	MSEXMTACONNS
Resource:	connection type1, connection type2, ...
Statistic Name:	
Inbound Bytes Total	<p>The total volume of message content received from this entity since the message transfer agent (MTA) initialization, measured in kilobytes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftExchg\Release2000\MTAConnectionsInboundAndOutboundTotalBytes.rpt</p>
Outbound Bytes Total	<p>The total volume of message content transmitted to this entity since the message transfer agent (MTA) initialization, measured in kilobytes</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftExchg\Release2000\MTAConnectionsInboundAndOutboundTotalBytes.rpt</p>

9.10.2. Microsoft Exchange 5.5

The following statistics can be collected by the System Activity Agent from Microsoft Exchange 5.5.

Parameter Hierarchy

Class:	MSExchangeIMC
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /MSExchange
TeamQuest Table Name:	MSExchangeIMC
Open Table Name:	MSEXIMC
Statistic Name:	
Inbound Bytes/Hr	The number of bytes transferred per hour to the Microsoft Exchange Server [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ IMCInboundAndOutboundBytesperHr.rpt
Inbound Messages/Hr	The number of messages transferred per hour to the Microsoft Exchange Server [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ IMCInboundAndOutboundMsgsperHr.rpt
Outbound Bytes/Hr	The number of bytes transferred per hour to the Internet [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ IMCInboundAndOutboundBytesperHr.rpt
Outbound Messages/Hr	The number of messages are transferred per hour to the Internet [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ IMCInboundAndOutboundMsgsperHr.rpt
Total Messages Queued	The total number of messages waiting in the Internet mail service queues [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ TotalRecipientsAndMsgsQueued.rpt
Total Recipients Queued	The total number of recipients in the Internet mail service queues [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ TotalRecipientsAndMsgsQueued.rpt

Class:	MSExchangeIS Private
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/systemname/MSExchange
TeamQuest Table Name:	MSExchangeIS Private
Open Table Name:	MSEXISPRIVATE
Statistic Name:	
Folder Opens/sec	The number of requests to open folders submitted per second to the information stores [Sequential = AVG Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ MsgsAndFoldersOpenedperSec.rpt
Message Opens/sec	The number of requests to open messages submitted per second to the information stores [Sequential = AVG Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ MsgsAndFoldersOpenedperSec.rpt
Messages Sent/min	The number of messages sent per minute to the other storage providers using the message transfer agent (MTA) [Sequential = AVG Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ MsgsSentAndSubmittedperMinute.rpt
Messages Submitted/min	The number of messages submitted per minute by clients [Sequential = AVG Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ MsgsSentAndSubmittedperMinute.rpt
Receive Queue Size	The number of messages in the private information stores receive queue [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ SendAndReceiveQueueSize.rpt
Send Queue Size	The number of messages in the private information stores send queue [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ SendAndReceiveQueueSize.rpt

Microsoft Windows Systems

Class:	MSExchangeMTA
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /MSExchange
TeamQuest Table Name:	MSExchangeMTA
Open Table Name:	MSEXMTA
Statistic Name:	
Disk File Reads/sec	The number of disk file read operations per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ MTADiskWritesAndReadsperSec.rpt
Disk File Writes/sec	The number of disk file write operations per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\MicrosoftExchg\Release55\ MTADiskWritesAndReadsperSec.rpt

9.10.3. Microsoft IIS Statistics

The following statistics can be collected by the System Activity Agent from the Microsoft Internet Information Server (IIS).

Parameter Hierarchy

Class:	Active Server Pages
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /IIS
TeamQuest Table Name:	Active Server Pages
Open Table Name:	ACTIVESERVPGS
Statistic Name:	
Sessions Total	The total number of sessions since the service was started [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftIIS\ ActServerPgsTotalSessions.rpt
Transactions Total	The total number of transactions since the service was started [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftIIS\ ActServerPgsTotalTrans.rpt

Class:	FTP Service
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /IIS
TeamQuest Table Name:	FTP Service
Open Table Name:	FTPSVC
Resource:	ftp service1, ftp service2, ...
Statistic Name:	
Total Connection Attempts (all instances)	<p>The number of connection attempts (from all instances) that have been made to the FTP service</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftIIS\FTPConnAtmpts.rpt</p>
Total Files Sent	<p>The number of files sent by the FTP service since startup</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftIIS\FTPFilesSent.rpt</p>
Class:	Internet Information Services Global
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /IIS
TeamQuest Table Name:	Internet Information Services Global
Open Table Name:	INTERNETINFORMATIONS
Statistic Name:	
File Cache Flushes	<p>The number of file cache flushes since server startup</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>\report\MicrosoftIIS\IISGlobalFileCacheFlushes.rpt</p>
File Cache Hits %	<p>The ratio of file cache hits to total cache requests</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>\report\MicrosoftIIS\IISGlobalFileCacheHitPercent.rpt</p>

Microsoft Windows Systems

Class:	Web Service
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /IIS
TeamQuest Table Name:	Web Service
Open Table Name:	WEBSVC
Resource:	web service1, web service2, ...
Statistic Name:	
Bytes Total/sec	<p>The sum of Bytes sent/sec and Bytes received/sec. This is the total number of bytes transferred by the Web service per second. [Sequential = AVG Non-Sequential = SUM] View Report: \report\MicrosoftIIS\ WebServiceBytesPerSec.rpt</p>
Maximum Connections	<p>The maximum number of simultaneous connections established with the Web service [Sequential = MAX Non-Sequential = MAX] View Report: \report\MicrosoftIIS\ WebServiceMaxConn.rpt</p>
Maximum ISAPI Extension Requests	<p>The maximum number of simultaneous connections established with the Web service [Sequential = MAX Non-Sequential = MAX] View Report: \report\MicrosoftIIS\ WebServiceMaxISAPIExtensionReqs.rpt</p>
Total Anonymous Users	<p>The total number of users who established an anonymous connection with the Web service (counted since service startup) [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftIIS\ WebServiceAnonymousUsers.rpt</p>
Total Connection Attempts (all instances)	<p>The number of connections that have been attempted using the Web service (counted since service startup) [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftIIS\ WebServiceConnAtmpts.rpt</p>
Total Not Found Errors	<p>The number of requests that could not be satisfied by the server because the requested document could not be found. These are generally reported as an HTTP 404 error code to the client. [Sequential = LST Non-Sequential = SUM] View Report: \report\MicrosoftIIS\ WebServiceNotFoundErrors.rpt</p>

9.10.4. Microsoft SQL Server

The following statistics can be collected by the System Activity Agent from the Microsoft SQL Server software and are created by a default SQL Server instance.

A named SQL Server instance has the Class name based on the SQL Server instance name. For example, a named instance of “SQL5” shows the “Buffer Manager” object as MSSQL\$SQL5.Buffer Manager.

The IT Resource name for the named SQL Server instance is also based on the SQL Server instance name. For our example, the IT Resource name would be /TeamQuest/System/systemname/SQL Server/SQL5.

Parameter Hierarchy

Class:	SQLServer:Access Methods
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/systemname/SQL Server/SQL
TeamQuest Table Name:	SQLServer:Access Methods
Open Table Name:	SQLSVRACCESSMETHODS
Statistic Name:	
Page Splits/sec	The number of page splits per second occurring as the result of index pages overflowing [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\SQL7\PageSplitsperSec.rpt
Table Lock Escalations/sec	The number of times locks on a table were escalated per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\TableLockEscalationsperSec.rpt
Workfiles Created/sec	The number of workfiles created in the last second [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\WorkFilesAndWorkTablesCreatedperSec.rpt
Worktables Created/sec	The number of worktables created in the last second [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\WorkFilesAndWorkTablesCreatedperSec.rpt

Microsoft Windows Systems

Class:	SQLServer:Buffer Manager
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /SQL Server/SQL
TeamQuest Table Name:	SQLServer:Buffer Manager
Open Table Name:	SQLSVRBUFFMGR
Statistic Name:	
Buffer Cache Hit Ratio	The percentage of pages that were found in the buffer cache without having to incur a read from disk [Sequential = AVG Non-Sequential = AVG] View Report: \report\SQLServer\SQL7\ BufferCacheHitRatio.rpt
Page Requests/sec	The number of requests for buffer page per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\SQL7\PageReqsperSec.rpt
Class:	SQLServer:Databases
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /SQL Server/SQL
TeamQuest Table Name:	SQLServer:Databases
Open Table Name:	SQLSVRDBS
Resource:	database1, database2, ...
Statistic Name:	
Repl. Trans. Rate	The replication transaction rate (per second) [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\ ReplicationTransperSec.rpt

Class:	SQLServer:General Statistics
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /SQL Server/SQL
TeamQuest Table Name:	SQLServer:General Statistics
Open Table Name:	SQLSVRGENERALSTATS
Statistic Name:	
Logins/sec	The total number of logins started per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\ UserLogins-LogoutsperSec.rpt
Logouts/sec	The total number of logouts started per second [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\ UserLogins-LogoutsperSec.rpt
Class:	SQLServer:Latches
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /SQL Server/SQL
TeamQuest Table Name:	SQLServer:Latches
Open Table Name:	SQLSVRLATCHES
Statistic Name:	
Latch Waits/sec	The number of latch requests per second that could not be granted immediately and had to wait before being granted [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\ SQLLatchWaitsperSec.rpt
Total Latch Wait Time (ms)	The total time in milliseconds (ms) that latch requests had to wait in the last second [Sequential = AVG Non-Sequential = SUM] View Report: \report\SQLServer\ SQLTotalLatchWaitTime.rpt

Microsoft Windows Systems

Class:	SQLServer:Locks
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /SQL Server/SQL
TeamQuest Table Name:	SQLServer:Locks
Open Table Name:	SQLSVRLOCKS
Resource:	lock type1, lock type2, ...
Statistic Name:	
Lock Waits/sec	<p>The number of lock requests per second that could not be satisfied immediately and required the caller to wait before being granted the lock</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\SQLServer\ TableLockWaitsperSec.rpt</p>
Class:	SQLServer:SQL Statistics
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /SQL Server/SQL
TeamQuest Table Name:	SQLServer:SQL Statistics
Open Table Name:	SQLSVRSQLSTATS
Statistic Name:	
SQL Compilations/sec	<p>The number of SQL compilations per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\SQLServer\ SQLCompileAndRecompileperSec.rpt</p>
SQL Re-Compilations/sec	<p>The number of SQL re-compilations per second</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: \report\SQLServer\ SQLCompileAndRecompileperSec.rpt</p>

Section 10

Network Applications

The Network Application Agent (**tqnpp**) passively monitors TCP/IP based applications (for example, web applications). It uses port mirroring on a local area network (LAN) switch to passively monitor TCP/IP traffic. The agent aggregates the monitored packets and calculates server and network response times, utilization, and throughput statistics for active applications. Optionally, the agent can calculate transaction rates and response times using an integrated set of heuristics.

This section contains a listing of the statistics collected by the agent:

- Performance Statistics (see 10.1)
- Client Statistics (see 10.2)
- Port Statistics (see 10.3)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

10.1. Performance Statistics

The Network Applications Agent (**tqnpp**) collects and stores data into the default aggregation set of the performance database. The statistics are classified by a hierarchy of key names. A statistic marked with an asterisk (*) is a derived statistic.

Parameter Hierarchy

Class:	Network Application
Subclass:	Agent
IT Resource Name:	/TeamQuest/Network/Traffic/ <i>systemname</i>
TeamQuest Table Name:	Network Application.Agent
Open Table Name:	NETAPPAGENT
Statistic Name:	
packet limit	Indicates whether or not the packet limit was reached [Sequential = LST Non-Sequential = LST]
packets dropped	The number of packets dropped during the sample interval. This value is not restricted by the list of ports for the agent. [Sequential = SUM Non-Sequential = SUM]
packets received	The number of packets received during the sample interval. This value is not restricted by the list of ports for the agent. [Sequential = SUM Non-Sequential = SUM]
packets recorded	The number of packets recorded for the sample interval [Sequential = SUM Non-Sequential = SUM]
port total	The number of ports being monitored [Sequential = SUM Non-Sequential = SUM]
processing limit	Indicates whether or not the processing time limit was reached [Sequential = LST Non-Sequential = LST]
processing time	The elapsed time in seconds to collect, process, and store the data [Sequential = SUM Non-Sequential = SUM]

Class:	Network Application
Subclass:	by Port
IT Resource Name:	/TeamQuest/Network/Traffic/ <i>systemname</i>
TeamQuest Table Name:	Network Application.by Port
Open Table Name:	NETAPPBYPOR
Resource:	port0, port1, port2, ...
Statistic Name:	
active clients	The number of active clients for the port [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/port/connection.rpt
connections	The number of new connections for the port [Sequential = SUM Non-Sequential = SUM] View Report: /report/network/application/port/connection.rpt
data segments/s*	The number of input and output data segments per second for the port
in Mbytes/s	The number of input bytes in megabytes per second for the port [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/port/transfers.rpt
in packets/s	The number of input packets per second for the port [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/port/packets.rpt
Mbytes/s*	The number of input and output bytes in megabytes per second for the port View Report: /report/network/application/port/sum-xfers.rpt
network time	The average amount of time it took a packet to traverse the network for the port [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/port/response.rpt
out Mbytes/s	The number of output bytes in megabytes per second for the port [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/port/transfers.rpt
out packets/s	The number of output packets per second for the port [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/port/packets.rpt
packets/s*	The number of input and output packets per second for the port View Report: /report/network/application/port/sum-pkts.rpt

Network Applications

reaction time	The average amount of time for the initial response to a request for the port [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/port/reaction.rpt
requests/s	The number of requests per second for the port [Sequential = AVG Non-Sequential = SUM]
response time*	The average amount of time for a request for the port View Report: /report/network/application/port/sum-resp.rpt
responses/s	The number of responses per second for the port [Sequential = AVG Non-Sequential = SUM]
server time	The average amount of time the server took to process a request for the port [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/port/response.rpt
terminations	The number of terminated connections for the port [Sequential = SUM Non-Sequential = SUM] View Report: /report/network/application/port/connection.rpt
transactions/s	The number of transactions per second for the port [Sequential = AVG Non-Sequential = SUM]
Class:	Network Application
Subclass:	Summary
IT Resource Name:	/TeamQuest/Network/Traffic/ <i>systemname</i>
TeamQuest Table Name:	Network Application.Summary
Open Table Name:	NETAPPSUM
Statistic Name:	
active clients	The number of active clients during the sample interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/network/application/summary/connection.rpt
active ports	The number of ports that had activity during the sample interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/network/application/summary/connection.rpt
connections	The number of new connections [Sequential = SUM Non-Sequential = SUM] View Report: /report/network/application/summary/connection.rpt
data segments/s*	The number of input and output data segments per second

in Mbytes/s	The number of input bytes in megabytes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/summary/transfers.rpt
in packets/s	The number of input packets per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/summary/packets.rpt
Mbytes/s*	The number of input and output bytes in megabytes per second View Report: /report/network/application/summary/sum-xfers.rpt
network time	The average amount of time it took a packet to traverse the network [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/summary/response.rpt
out Mbytes/s	The number of output bytes in megabytes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/summary/transfers.rpt
out packets/s	The number of output packets per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/summary/packets.rpt
packets/s*	The number of input and output packets per second View Report: /report/network/application/summary/sum-pkts.rpt
reaction time	The average amount of time for the initial response to a request [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/summary/reaction.rpt
requests/s	The number of requests per second [Sequential = AVG Non-Sequential = SUM]
response time*	The average amount of time for a request View Report: /report/network/application/summary/sum-resp.rpt
responses/s	The number of responses per second [Sequential = AVG Non-Sequential = SUM]
server time	The average amount of time the server took to process a request [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/application/summary/response.rpt
terminations	The number of terminated connections [Sequential = SUM Non-Sequential = SUM] View Report: /report/network/application/summary/connection.rpt

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.AgentInterval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

10.2. Client Statistics

The Network Application Agent stores detailed client information in the Network Application:Client table in the TeamQuest performance database. The values of the statistics represent a summary of all of the activity for each client system that had a conversation with the configured system.

Note: This data is available only when the Client Table setting is configured to ON.

Table Field Hierarchy

Class:	Network Application
Subclass:	Client
IT Resource Name:	/TeamQuest/Network/Traffic/systemname
TeamQuest Table Name:	Network Application.Client
Open Table Name:	NETAPPCLI
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Client	The client system name [Sequential = ID Non-Sequential = ID]
Connections	The number of new connections for the client [Sequential = SUM Non-Sequential = SUM]
Data Segments	The number of input and output data segments for the client [Sequential = SUM Non-Sequential = SUM]
Interval	The expected data sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
Mbytes	The number of input and output bytes in megabytes for the client [Sequential = SUM Non-Sequential = SUM]
Mbytes In	The number of input bytes in megabytes for the client [Sequential = SUM Non-Sequential = SUM]
Mbytes Out	The number output bytes in megabytes for the client [Sequential = SUM Non-Sequential = SUM]
Network Time	The average amount of time it took a packet to traverse the network for the client [Sequential = AVG Non-Sequential = AVG]

Network Applications

Packets	The number of input and output packets for the client [Sequential = SUM Non-Sequential = SUM]
Packets In	The number of input packets for the client [Sequential = SUM Non-Sequential = SUM]
Packets Out	The number of output packets for the client [Sequential = SUM Non-Sequential = SUM]
Ports	This statistic is not available. The value is displayed as <N/A>.
Reaction Time	The average amount of time for the initial response to a request for the client [Sequential = AVG Non-Sequential = AVG]
Requests	The number of requests to the client [Sequential = SUM Non-Sequential = SUM]
Responses	The number of responses for the client [Sequential = SUM Non-Sequential = SUM]
Server Time	The average amount of time the server took to process a request for the client [Sequential = AVG Non-Sequential = AVG]
System	The server system name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Terminations	The number of terminated connections for the client [Sequential = SUM Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Transaction Time	The average transaction response time for the client. This statistic will be set to <N/A> if the Record Transactions setting is set to OFF. [Sequential = AVG Non-Sequential = AVG]
Transactions	The number of transactions for the client. This statistic will be set to <N/A> if the Record Transactions setting is set to OFF. [Sequential = SUM Non-Sequential = SUM]

10.3. Port Statistics

The Network Application Agent stores detailed port information in the Network Application:Port table in the TeamQuest performance database. The values of the statistics represent a summary of all of the activity for each port system for the configured system.

Table Field Hierarchy

Class:	Network Application
Subclass:	Port
IT Resource Name:	/TeamQuest/Network/Traffic/ <i>systemname</i>
TeamQuest Table Name:	Network Application.Port
Open Table Name:	NETAPPPORT
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Clients	The number of active clients for the port [Sequential = SUM Non-Sequential = SUM]
Connections	The number of new connections for the port [Sequential = SUM Non-Sequential = SUM]
Data Segments	The number of input and output data segments for the port [Sequential = SUM Non-Sequential = SUM]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Mbytes	The number of input and output bytes in megabytes for the port [Sequential = SUM Non-Sequential = SUM]
Mbytes In	The number of input bytes in megabytes for the port [Sequential = SUM Non-Sequential = SUM]
Mbytes Out	The number output bytes in megabytes for the port [Sequential = SUM Non-Sequential = SUM]
Network Time	The average amount of time it took a packet to traverse the network for the port [Sequential = AVG Non-Sequential = AVG]
Packets	The number of input and output packets for the port [Sequential = SUM Non-Sequential = SUM]
Packets In	The number of input packets for the port [Sequential = SUM Non-Sequential = SUM]

Network Applications

Packets Out	The number of output packets for the port [Sequential = SUM Non-Sequential = SUM]
Port	The port name [Sequential = ID Non-Sequential = ID]
Reaction Time	The average amount of time for the initial response to a request for the port [Sequential = AVG Non-Sequential = AVG]
Requests	The number of requests for the port [Sequential = SUM Non-Sequential = SUM]
Responses	The number of responses for the port [Sequential = SUM Non-Sequential = SUM]
Server Time	The average amount of time the server took to process a request for the port [Sequential = AVG Non-Sequential = AVG]
System	The server system name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Terminations	The number of terminated connections for the port [Sequential = SUM Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Transaction Time	The average transaction response time for the port. This statistic will be set to <N/A> if the Record Transactions setting is set to OFF. [Sequential = AVG Non-Sequential = AVG]
Transactions	The number of transactions for the port. This statistic will be set to <N/A> if the Record Transactions setting is set to OFF. [Sequential = SUM Non-Sequential = SUM]

Section 11

Network Devices

The Network Device Agent (**tqndp**) monitors Simple Network Management Protocol (SNMP) enabled devices for performance data. Those devices include, but are not limited to, switches and routers. The agent collects utilization, throughput, and error-related performance data.

This section contains a listing of the statistics collected by the agent:

- Summary Statistics (see 11.1)
- Interface Statistics (see 11.2)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

11.1. Summary Statistics

The Network Device Agent (**tqndp**) collects and stores data into the default aggregation set of the performance database. The statistics are classified by a hierarchy of key names. The values of the following statistics represent a summary of all the activity for each configured network device. A statistic marked with an asterisk (*) is a derived statistic.

Parameter Hierarchy

Class:	Network Device
Subclass:	Summary
IT Resource Name:	/TeamQuest/Network/Devices/ <i>systemname</i>
TeamQuest Table Name:	Network Device.Summary
Open Table Name:	NETDEVSUM
Statistic Name:	
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgBusy1	The 1-minute exponentially-decayed moving average of the CPU busy percentage. This statistic is available only for selected Cisco devices. [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/cisco-cpu.rpt
avgBusy5	The 5-minute exponentially-decayed moving average of the CPU busy percentage. This statistic is available only for selected Cisco devices. [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/cisco-cpu.rpt
cpmCPUTotal1min	The overall percentage of CPU busy in the last 1-minute sampling interval [Sequential = LST Non-Sequential = LST] View Report: /report/network/device/cisco-cpu.rpt
cpmCPUTotal1minRev	The overall percentage of CPU busy in the last 1-minute sampling interval [Sequential = LST Non-Sequential = LST] View Report: /report/network/device/cisco-cpu.rpt
cpmCPUTotal5min	The overall percentage of CPU busy in the last 5-minute sampling interval [Sequential = LST Non-Sequential = LST] View Report: /report/network/device/cisco-cpu.rpt

cpmCPU5minRev	The overall percentage of CPU busy in the last 5-minute sampling interval [Sequential = LST Non-Sequential = LST] View Report: /report/network/device/cisco-cpu.rpt
discards/s*	The number of input and output discards per second
errors/s*	The number of input and output errors per second
in discards/s*	The number of input discards per second
in errors/s*	The number of input errors per second
in Kbytes/s*	The number of input bytes per second in kilobytes
in packets/s*	The number of input packets per second
Kbytes/s*	The number of input and output bytes per second in kilobytes
out discards/s*	The number of output discards per second
out errors/s*	The number of output errors per second
out Kbytes/s*	The number of output bytes per second in kilobytes
out packets/s*	The number of output packets per second
packets/s*	The number of input and output packets per second View Report: /report/network/device/sum-pkts.rpt
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.AgentInterval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

11.2. Interface Statistics

The Network Device Agent (**tqndp**) collects and stores data into the default aggregation set of the performance database. The statistics are classified by a hierarchy of key names. The values of the following statistics represent a summary of all the activity for each interface of the configured network device. A statistic marked with an asterisk (*) is a derived statistic.

Parameter Hierarchy

Class:	Network Device
Subclass:	by Interface
IT Resource Name:	/TeamQuest/Network/Devices/ <i>systemname</i>
TeamQuest Table Name:	Network Device.by Interface
Open Table Name:	NETDEVBYINTERFACE
Resource:	interface0, interface1, interface2, ...
Statistic Name:	
discards/s*	The number of input and output discards per second
errors/s*	The number of input and output errors per second
ifspeed	An estimate of the interface's current bandwidth in bits per second. For interfaces that do not vary in bandwidth or for those where no accurate estimation can be made, this value may represent the nominal bandwidth. [Sequential = AVG Non-Sequential = SUM]
in bcpackets/s	The number of packets addressed to a broadcast address at this sub-layer [Sequential = AVG Non-Sequential = SUM]
in discards/s	The number of input discards per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/discards.rpt
in errors	The number of input errors per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/errors.rpt
in Kbytes/s	The number of input bytes per second in kilobytes [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/transfers.rpt
in mcpackets/s	The number of packets addressed to a multicast address at this sub-layer. For a MAC layer protocol, this includes the Group and Functional addresses. [Sequential = AVG Non-Sequential = SUM]
in nupackets/s	The number of input non-unicast packets per second. This field contains the sum of mcpackets and bcpackets if the interface supports bcpackets and mcpackets. [Sequential = AVG Non-Sequential = SUM]

in packets/s*	The number of input packets per second View Report: /report/network/device/packets.rpt
in upackets/s	The number of input unicast packets per second [Sequential = AVG Non-Sequential = SUM]
Kbytes/s*	The number of input and output bytes per second in kilobytes View Report: /report/network/device/sum-xfers.rpt
out bcpackets/s	The total number of packets that higher-level protocols requested to be transmitted and that were addressed to a broadcast address at this sub-layer, including the packets that were discarded or not sent [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/sum-pkts.rpt
out discards/s	The number of output discards per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/discards.rpt
out errors/s	The number of output errors per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/errors.rpt
out Kbytes/s	The number of output bytes per second in kilobytes [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/transfers.rpt
out mcpackets/s	The total number of packets that higher-level protocols requested to be transmitted and that were addressed to a multicast address at this sub-layer, including the packets that were discarded or not sent. For a MAC layer protocol, this includes Group and Functional addresses. [Sequential = AVG Non-Sequential = SUM] View Report: /report/network/device/sum-pkts.rpt
out nupackets/s	The number of output non-unicast packets per second. This field contains the sum of mcpackets and bcpackets if the interface supports bcpackets and mcpackets. [Sequential = AVG Non-Sequential = SUM]
out packets/s*	The number of output packets per second View Report: /report/network/device/packets.rpt
out upackets/s	The number of output unicast packets per second [Sequential = AVG Non-Sequential = SUM]
packets/s*	The number of input and output packets per second

Section 12

Oracle Database Server

The Oracle Data Agent (**tqorap**) collects information on Oracle instances. The agent obtains instance summary data pertaining to memory, disk I/O, system global area (SGA), network, rollback segments, block contention, and latches. Detailed metrics about sessions, datafiles, rollback segments, library cache, row cache, block contention, latches, system wait events, session wait events, top SQL cursors, and instance configuration information are also stored in the performance database.

The Oracle Alarm Agent (**tqoraalm**) monitors conditions for Oracle instances and Oracle listeners, and generates alarms when appropriate. The agent observes alarm conditions for the Oracle instance status, Oracle listener status, errors detected in the alert.log file, locks held, and database object capacity.

This section contains a listing of the statistics collected by the agents:

- Alert Log Alarm Statistics (see 12.1)
- Block Contention Wait Statistics (see 12.2)
- Datafile Capacity Statistics (see 12.3)
- Datafile I/O Statistics (see 12.4)
- Instance Alarm Statistics (see 12.5)
- Instance Data Statistics (see 12.6)
- Latch Statistics (see 12.7)
- Library Cache Statistics (see 12.8)
- Listener Alarm Statistics (see 12.9)
- Lock Alarm Statistics (see 12.10)
- Rollback Segment Statistics (see 12.11)
- Row Cache Statistics (see 12.12)
- Segment Alarm Statistics (see 12.13)
- Session Statistics (see 12.14)
- Session Wait Statistics (see 12.15)
- System Parameters Statistics (see 12.16)
- System Statistics (see 12.17)
- System Wait Event Statistics (see 12.18)
- Top SQL Cursors Statistics (see 12.19)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

12.1. Alert Log Alarm Statistics

The Oracle.Alarm_AlertLog table stores alarm information about the error codes within the alert.log file.

Table Field Hierarchy

Class:	Oracle
Subclass:	Alarm_AlertLog
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Alarm_AlertLog
Open Table Name:	ORAALARMALERTLOG
Collection interval:	N/A
Default retention:	30 days
Table type:	Event

Statistic Name	Description
Alarm_ID	The user-defined alarm identifier assigned to the alarm [Non-Sequential = ID]
error_code	The error code encountered in the alert.log file [Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Non-Sequential = ID]
Match_ID	A tab-delimited concatenation of the time consolidation fields for the record. This is a hidden field and is for internal use only. [Non-Sequential = ID]
Sequence	A sequential number assigned to the alarm for uniqueness [Non-Sequential = ID]
Severity	The severity of the alarm (Normal, Warning, Minor, Major, or Critical) [Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

12.2. Block Contention Wait Statistics

The Oracle.Waitstat table stores detailed information about block contention within the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	Waitstat
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Waitstat
Open Table Name:	ORAWAITSTAT
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avg_waittime	The average amount of time per wait by the operation for the class of block during the interval. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = AVG Non-Sequential = AVG]
class	The class of block [Sequential = ID Non-Sequential = ID]
count	The number of waits by the operation for the class of block during the interval [Sequential = SUM Non-Sequential = SUM]
count_t	The total number of waits by the operation for the class of block since the instance started [Sequential = LST Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
waittime	The amount of wait time in seconds by the operation for the class of block during the interval. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
waittime_t	The total amount of wait time in seconds for the waits by the operation for the class of block since the instance started. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = LST Non-Sequential = SUM]

12.3. Datafile Capacity Statistics

The Oracle.Datafile_Capacity table stores detailed capacity information about the datafiles used by the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	Datafile_Capacity
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Datafile_Capacity
Open Table Name:	ORADATAFILECAPACITY
Collection interval:	10 minutes (default)
Default retention:	10 days
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avg_frag	The average size in blocks of a free extent in the datafile [Sequential = AVG Non-Sequential = AVG]
create_MB	The size in megabytes (MB) of the datafile when it was initially created [Sequential = LST Non-Sequential = SUM]

file_id	The file identification number [Sequential = ID Non-Sequential = ID]
filename	The name of the datafile. Up to 256 characters are displayed. [Sequential = ID Non-Sequential = ID]
fragments	The number of free extents in the datafile [Sequential = SUM Non-Sequential = SUM]
free_blks	The current amount of free space in blocks within the datafile [Sequential = LST Non-Sequential = SUM]
free_MB	The current amount of free space in megabytes (MB) within the datafile [Sequential = LST Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
max_frag	The size in blocks of the largest free extent in the datafile [Sequential = MAX Non-Sequential = MAX]
min_frag	The size in blocks of the smallest free extent in the datafile [Sequential = MIN Non-Sequential = MIN]
pct_free	The percentage of the datafile that is free [Sequential = LST Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
status	The status of the datafile [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
tablespace_name	The name of the tablespace to which the file belongs [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
total_blks	The current size in blocks of the datafile. The value will be zero if the file is inaccessible. [Sequential = LST Non-Sequential = SUM]
total_MB	The current size in megabytes of the datafile. The value will be zero (0) if the file is inaccessible. [Sequential = LST Non-Sequential = SUM]

12.4. Datafile I/O Statistics

The Oracle.Datafile_IO table stores detailed I/O information about the datafiles used by the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	Datafile_IO
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Datafile_IO
Open Table Name:	ORADATAFILEIO
Collection interval:	1 minute (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avg_readtime	The average amount of time in seconds spent doing a read from the datafile during the interval. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = AVG Non-Sequential = AVG]
avg_writetime	The average amount of time in seconds spent doing a write to the datafile during the interval. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = AVG Non-Sequential = AVG]
file_id	The file identification number [Sequential = ID Non-Sequential = ID]
filename	The name of the datafile. Up to 256 characters are displayed. [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
phys_blk_rds	The number of physical blocks read from the datafile during the interval. The value reflects requests to the operating system, which may not directly correspond to a physical block read from the disk on which the datafile resides. [Sequential = SUM Non-Sequential = SUM]

phys_blk_rds_t	<p>The total number of physical blocks read from the datafile since the instance was started. The value reflects requests to the operating system and may not directly correspond to a physical block read from the disk on which the datafile resides.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
phys_blk_wrts	<p>The number of physical blocks written to the datafile during the interval. The value reflects requests to the operating system, which may not directly correspond to a physical block written to the disk on which the datafile resides.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
phys_blk_wrts_t	<p>The total number of physical blocks written to the datafile since the instance was started. The value reflects requests to the operating system, which may not directly correspond to a physical block written to the disk on which the datafile resides.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
phys_rds	<p>The number of physical reads from the datafile during the interval. The value reflects requests to the operating system, which may not directly correspond to a physical read from the disk on which the datafile resides.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
phys_rds_t	<p>The total number of physical reads from the datafile since the instance was started. The value reflects requests to the operating system, which may not directly correspond to a physical read from the disk on which the datafile resides.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
phys_wrts	<p>The number of physical writes to the datafile during the interval. The value reflects requests to the operating system, which may not directly correspond to a physical write to the disk on which the datafile resides.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
phys_wrts_t	<p>The total number of physical writes to the datafile since the instance was started. The value reflects requests to the operating system, which may not directly correspond to a physical write to the disk on which the datafile resides.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
readtime	<p>The amount of time in seconds spent doing reads from the datafile during the interval. This value is zero when the Oracle <code>timed_statistics</code> initialization parameter is set to <code>FALSE</code>.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
readtime_t	<p>The total amount of time in seconds spent doing reads from the datafile since the instance started. This value is zero when the Oracle <code>timed_statistics</code> initialization parameter is set to <code>FALSE</code>.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
status	<p>The status of the datafile</p> <p>[Sequential = ID Non-Sequential = ID]</p>

System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
tablespace_name	The name of the tablespace to which the file belongs [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
writetime	The amount of time in seconds spent doing writes to the datafile during the interval. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
writetime_t	The total amount of time in seconds spent doing writes to the datafile since the instance started. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = LST Non-Sequential = SUM]

12.5. Instance Alarm Statistics

The Oracle.Alarm_Instance table stores alarm information about the status of Oracle instances.

Table Field Hierarchy

Class:	Oracle
Subclass:	Alarm_Instance
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Alarm_Instance
Open Table Name:	ORAALARMINSTANCE
Collection interval:	60 seconds (default)
Default retention:	30 days
Table type:	Performance

Statistic Name	Description
Alarm_ID	The user-defined alarm identifier assigned to the alarm [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Match_ID	A tab-delimited concatenation of the time consolidation fields for the record. This is a hidden field and is for internal use only. [Sequential = ID Non-Sequential = ID]

Severity	The severity of the alarm (Normal, Warning, Minor, Major, or Critical) [Sequential = LST Non-Sequential = ID]
status	The status of the instance (Up or Down) [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

12.6. Instance Data Statistics

This subsection contains a list of the Oracle statistics collected by the Oracle Data Agent (**tqorap**). A statistic marked with an asterisk (*) is a derived statistic.

Parameter Hierarchy

Class:	Oracle
Subclass:	DBWR
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Oracle/ <i>instancename</i>
TeamQuest Table Name:	Oracle.DBWR
Open Table Name:	ORADBWR
Resource:	instance1, instance2, ...
Statistic Name:	
BuffersScanned	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
Chkpoints	The number of times the database writer (DBWR) was asked to scan the cache and write all blocks marked for a checkpoint or the end of recovery. This statistic is always larger than background checkpoints completed. [Sequential = SUM Non-Sequential = SUM]
CrossInstancewrites	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
Freebuffersfound	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
LruScans	The number of times that the database writer (DBWR) scans the least recently used (LRU) queue looking for buffers to write. This count includes scans to fill a batch being written for another purpose (such as a checkpoint). [Sequential = SUM Non-Sequential = SUM]

MakeFreeRequests	The number of requests to the database writer (DBWR) to make some free buffers for the least recently used (LRU). This statistic is not available with Oracle 11.2 and later. [Sequential = SUM Non-Sequential = SUM]
RevisitedWriteBuffer	The number of times that the database writer (DBWR) tried to save a buffer for writing and found that it was already in the write batch. This statistic measures the amount of useless work that the DBWR had to do in trying to fill the batch. Many sources contribute to a write batch. If the same buffer from different sources is considered for adding to the write batch, then all but the first attempt will be useless because the buffer is already marked as being written. [Sequential = SUM Non-Sequential = SUM]
SummedScanDepth	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
TransTableWrites	The number of rollback segment headers written by the database writer (DBWR). This statistic indicates how many hot buffers were written, causing a user process to wait while the write completed. [Sequential = SUM Non-Sequential = SUM]
UndoBlockWrites	The number of rollback segment blocks written by the database writer (DBWR) [Sequential = SUM Non-Sequential = SUM]
Class:	Oracle
Subclass:	Disk
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Disk
Open Table Name:	ORADISK
Resource:	instance1, instance2, ...
Statistic Name:	
AvgDiskFree	The average size of allocated disk space not being used in bytes during the interval [Sequential = AVG Non-Sequential = SUM]
DiskAllocated	The size of allocated disk space in bytes [Sequential = LST Non-Sequential = SUM]
DiskFree	The size of allocated disk space not being used in bytes [Sequential = LST Non-Sequential = SUM]
FreeDiskSpace MB*	The amount of allocated disk space that is not being used in megabytes (MB) View Report: /report/oracle/DiskUsage.rpt
MaxDiskFree	The maximum size of allocated disk space not being used in bytes during the interval [Sequential = MAX Non-Sequential = SUM]

PhysicalBlockReads	The number of physical blocks read. The value reflects requests to the operating system, which may not directly correspond to a physical block read from a disk. [Sequential = SUM Non-Sequential = SUM]
PhysicalBlockWrites	The number of physical blocks written. The value reflects requests to the operating system, which may not directly correspond to a physical block written to a disk. [Sequential = SUM Non-Sequential = SUM]
PhysicalReads	The number of physical reads. The value reflects requests to the operating system, which may not directly correspond to a physical read to a disk. [Sequential = SUM Non-Sequential = SUM]
PhysicalReadsDirect	The number of reads directly from disk, bypassing the buffer cache. For example, in high bandwidth, data-intensive operations such as parallel query, reads of disk blocks bypass the buffer cache to maximize transfer rates and to prevent the premature aging of shared data blocks resident in the buffer cache. [Sequential = SUM Non-Sequential = SUM]
PhysicalReadsDirect(LOB)	The number of large object reads directly from disk, bypassing the buffer cache [Sequential = SUM Non-Sequential = SUM]
PhysicalWrites	The number of physical writes. The value reflects requests to the operating system, which may not directly correspond to a physical write to a disk. [Sequential = SUM Non-Sequential = SUM]
PhysicalWritesDirect	The number of writes directly to disk, bypassing the buffer cache (as in a direct load operation) [Sequential = SUM Non-Sequential = SUM]
PhysicalWritesDirect(LOB)	The number of large object writes directly to the disk, bypassing the buffer cache [Sequential = SUM Non-Sequential = SUM]
PhysicalWritesNonCheckpoint	The number of times a buffer is written for reasons other than advancement of the checkpoint. Used as a metric for determining the I/O overhead imposed by setting the FAST_START_IO_TARGET parameter to limit recovery I/Os. Essentially this statistic measures the number of writes that would have occurred had there been no checkpointing. Subtracting this value from physical writes gives the extra I/O for checkpointing. [Sequential = SUM Non-Sequential = SUM]

Pinned_Buffers_Inspected	The number of times a user process, when scanning the tail of the replacement list looking for a buffer to reuse, encountered a cold buffer that was pinned or had a waiter that was about to pin it. This occurrence is uncommon, because a cold buffer should not be pinned very often. [Sequential = SUM Non-Sequential = SUM]
TotalDiskSpace MB*	The amount of disk space allocated in megabytes (MB) View Report: /report/oracle/DiskUsage.rpt
Class:	Oracle
Subclass:	Enqueue
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Enqueue
Open Table Name:	ORAENQ
Resource:	instance1, instance2, ...
Statistic Name:	
Conversions	The total number of conversions of the state of table or row lock [Sequential = SUM Non-Sequential = SUM]
Deadlocks	The total number of deadlocks between table or row locks in different sessions [Sequential = SUM Non-Sequential = SUM]
Releases	The total number of table or row locks released [Sequential = SUM Non-Sequential = SUM]
Requests	The total number of table or row locks acquired [Sequential = SUM Non-Sequential = SUM]
Timeouts	The total number of table and row locks (acquired and converted) that timed out before completion [Sequential = SUM Non-Sequential = SUM]
Waits	The total number of waits that occurred during an enqueue convert or get because the enqueue get was deferred [Sequential = SUM Non-Sequential = SUM]

Oracle Database Server

Class:	Oracle
Subclass:	GlobalCache
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.GlobalCache
Open Table Name:	ORAGLOBALCACHE
Resource:	instance1, instance2, ...
Statistic Name:	
BlocksCorrupt	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
Converts	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
ConvertTime	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
ConvertTimeouts	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
CrBlockBuildTime	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
CrBlockFlushTime	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
CrBlockRecvTime	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
CrBlockSendTime	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
CrBlocksReceived	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
CrBlocksServed	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
Defers	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
FreelistWaits	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]

Gets	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
GetTime	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
PrepareFailures	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
Class:	Oracle
Subclass:	GlobalLock
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.GlobalLock
Open Table Name:	ORAGLOBALLOCK
Resource:	instance1, instance2, ...
Statistic Name:	
Async_Converts	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
Async_Gets	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
ConvertTime	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
GetTime	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
Releases	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
SyncConverts	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
SyncGets	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]

Class:	Oracle
Subclass:	Latch Summary
IT Resource Name:	<i>/TeamQuest/System/systemname/Oracle/instancename</i>
TeamQuest Table Name:	Oracle.Latch Summary
Open Table Name:	ORALATCHSUM
Resource:	instance1, instance2, ...
Statistic Name:	
LatchGets	The number of times a latch was obtained with a wait [Sequential = SUM Non-Sequential = SUM]
LatchImmediateGets	The number of times a latch was obtained without a wait [Sequential = SUM Non-Sequential = SUM]
LatchImmediateMisses	The number of times a latch failed to be obtained without a wait [Sequential = SUM Non-Sequential = SUM]
LatchMisses	The number of times a latch was obtained with a wait but failed on the first try [Sequential = SUM Non-Sequential = SUM]
Class:	Oracle
Subclass:	Library
IT Resource Name:	<i>/TeamQuest/System/systemname/Oracle/instancename</i>
TeamQuest Table Name:	Oracle.Library
Open Table Name:	ORALIB
Resource:	instance1, instance2, ...
Statistic Name:	
CacheGets	The number of lock requests for library cache objects [Sequential = SUM Non-Sequential = SUM]
CacheGetHits	The number of requested objects that were already allocated in the library cache [Sequential = SUM Non-Sequential = SUM]
CachePins	The number of pin requests for the library cache objects [Sequential = SUM Non-Sequential = SUM]
CachePinHits	The number of pin requested objects that were already allocated in the library cache [Sequential = SUM Non-Sequential = SUM]
Reloads	The number of pin requests for library cache objects that require reloading [Sequential = SUM Non-Sequential = SUM]

Class:	Oracle
Subclass:	Memory
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Oracle/ <i>instancename</i>
TeamQuest Table Name:	Oracle.Memory
Open Table Name:	ORAMEM
Resource:	instance1, instance2, ...
Statistic Name:	
AvgSessionMemory	The average amount of session memory allocated in bytes during the interval [Sequential = AVG Non-Sequential = SUM]
AvgSgaFreeMemory	The average amount of system global area (SGA) memory not being used in bytes during the interval [Sequential = AVG Non-Sequential = SUM]
MaxSgaFreeMemory	The maximum amount of SGA memory not being used in bytes during the interval [Sequential = MAX Non-Sequential = MAX]
SessionMaxMemory	The maximum amount of session memory allocated for existing sessions in bytes [Sequential = LST Non-Sequential = SUM]
SessionMax Memory MB*	The maximum amount of session memory allocated for existing sessions in megabytes (MB) View Report: /report/oracle/SessionMem.rpt
SessionMemory	The amount of session memory allocated in bytes [Sequential = LST Non-Sequential = SUM]
SessionMemory MB*	The amount of session memory allocated in megabytes (MB) View Report: /report/oracle/SessionMem.rpt
SgaFreeMemory	The amount of SGA memory not being used in bytes [Sequential = LST Non-Sequential = SUM]
SgaFreeMemory MB*	The amount of SGA memory not being used in megabytes (MB) View Report: /report/oracle/SgaMem.rpt
SgaTotalMemory	The size of SGA memory in bytes [Sequential = LST Non-Sequential = SUM]
SgaTotalMemory MB*	The size of SGA memory in megabytes (MB) View Report: /report/oracle/SgaMem.rpt

Class:	Oracle
Subclass:	Network
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Network
Open Table Name:	ORANET
Resource:	instance1, instance2, ...
Statistic Name:	
BckgrndChkpnt Completed	The number of checkpoints completed by the background process. This statistic is incremented when the background process successfully advances the thread checkpoint. [Sequential = SUM Non-Sequential = SUM]
BckgrndChkpntStarted	The number of checkpoints started by the background process. This statistic can be larger than background checkpoints completed if a new checkpoint overrides an incomplete checkpoint or if a checkpoint is currently under way. This statistic includes only checkpoints of the redo thread. [Sequential = SUM Non-Sequential = SUM]
BckgrndTimeouts	The number of timeouts by the background process [Sequential = SUM Non-Sequential = SUM]
BranchNodeSplits	The number of times an index branch block was split because of the insertion of an additional value [Sequential = SUM Non-Sequential = SUM]
BufferIsNotPinnedCnt	The number of times a buffer was free when visited. This statistic is useful only for internal debugging purposes. [Sequential = SUM Non-Sequential = SUM]
BufferIsPinnedCnt	The number of times a buffer was pinned when visited. This statistic is useful only for internal debugging purposes. [Sequential = SUM Non-Sequential = SUM]
ClientBytesReceived	The number of bytes received from a client using Oracle Networking [Sequential = SUM Non-Sequential = SUM]
ClientBytesReceived/s*	The number of bytes per second received from a client using Oracle Networking View Report: /report/oracle/NetTransfers.rpt
ClientBytesSent	The number of bytes sent to a client using Oracle Networking [Sequential = SUM Non-Sequential = SUM]
ClientBytesSent/s*	The number of bytes per second sent to a client using Oracle Networking View Report: /report/oracle/NetTransfer.rpt
ClientRoundtrips	The number of Oracle Networking messages sent to and received from the client [Sequential = SUM Non-Sequential = SUM]

ClientRoundtrips/s*	The number of Oracle Networking messages per second sent to and received from the client View Report: /report/oracle/NetRoundtrips.rpt
DblinkBytesReceived	The number of bytes received from a database link using Oracle Networking [Sequential = SUM Non-Sequential = SUM]
DblinkBytesReceived/s*	The number of bytes per second received from a database link using Oracle Networking View Report: /report/oracle/NetTransfer.rpt
DblinkBytesSent	The number of bytes sent to a database link using Oracle Networking [Sequential = SUM Non-Sequential = SUM]
DblinkBytesSent/s*	The number of bytes per second sent to a database link using Oracle Networking View Report: /report/oracle/NetTransfers.rpt
DblinkRoundtrips	The number of Oracle Networking messages sent over and received from a database link [Sequential = SUM Non-Sequential = SUM]
DblinkRoundtrips/s*	The number Oracle Networking messages per second sent over and received from a database link View Report: /report/oracle/NetRoundtrips.rpt
UnnecessaryclnupSCN	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
Class:	Oracle
Subclass:	ParallelOps
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.ParallelOps
Open Table Name:	ORAPARALLELOPS
Resource:	instance1, instance2, ...
Statistic Name:	
Downgraded1to25pct	The number of times parallel execution was requested and the degree of parallelism was reduced because of insufficient parallel execution servers [Sequential = SUM Non-Sequential = SUM]
Downgraded25to50pct	The number of times parallel execution was requested and the degree of parallelism was reduced because of insufficient parallel execution servers [Sequential = SUM Non-Sequential = SUM]

Downgraded50to75pct	The number of times parallel execution was requested and the degree of parallelism was reduced because of insufficient parallel execution servers [Sequential = SUM Non-Sequential = SUM]
Downgraded75to99pct	The number of times parallel execution was requested and the degree of parallelism was reduced because of insufficient parallel execution servers [Sequential = SUM Non-Sequential = SUM]
DowngradedtoSerial	The number of times parallel execution was requested but execution was serial because of insufficient parallel execution servers [Sequential = SUM Non-Sequential = SUM]
NotDowngraded	The number of times parallel execution was executed at the requested degree of parallelism [Sequential = SUM Non-Sequential = SUM]

Note: A statistic marked with an asterisk (*) is a derived statistic and therefore is only available for viewing in TeamQuest Analyzer:

Class:	Oracle
Subclass:	Ratio
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
Resource:	instance1, instance2, ...
Statistic Name:	
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
BlockChangesPer Transaction	The amount of data manipulation language (DML) work performed by each transaction View Report: /report/oracle/BlockRate.rpt
BlockGetRate	The rate at which the application system references the database View Report: /report/oracle/BlockRate.rpt
CacheHitRatio	The buffer cache hit ratio View Report: /report/oracle/BufferCache.rpt
CallRate	The work demand rate being placed on the instance from all work sources View Report: /report/oracle/CallRate.rpt

ChangedBlockRatio	The balance between queries and data manipulation language (DML) View Report: /report/oracle/BlockRate.rpt
ConsistentChangeRatio	The extent to which applications are having to exercise the read consistency mechanism
ConsistentChanges*	The number of consistent get operations that cannot accept the block in its current state [Sequential = SUM Non-Sequential = SUM]
ConsistentGets*	The number of requests to the buffer manager to locate a database block as part of a consistent-read operation [Sequential = SUM Non-Sequential = SUM]
ContinuedRowRatio	The percentage of total rows retrieved that were continued rows View Report: /report/oracle/RowRatio.rpt
DbBlockChanges*	The number of database blocks in memory that created a “dirty block” [Sequential = SUM Non-Sequential = SUM]
DbBlockGets*	The number of requests to the buffer manager for a database block regardless of read consistency [Sequential = SUM Non-Sequential = SUM]
DiskSorts*	The number of sorts that allocated work space on disk [Sequential = SUM Non-Sequential = SUM]
MemorySorts*	The number of sorts that did not require allocation of work space on disk [Sequential = SUM Non-Sequential = SUM]
ParseCount*	The number of parse calls received from connected applications and recursive calls [Sequential = SUM Non-Sequential = SUM]
PhysicalBlockReads	The number of physical blocks read. The value reflects requests to the operating system, which may not directly correspond to a physical block read from a disk. [Sequential = SUM Non-Sequential = SUM]
PhysicalBlockReads/s	The number of physical blocks read per second View Report: /report/oracle/PhysBlockIO.rpt
PhysicalBlockWrites	The number of physical blocks written. The value reflects requests to the operating system, which may not directly correspond to a physical block written to a disk. [Sequential = SUM Non-Sequential = SUM]
PhysicalBlockWrites/s	The number of physical blocks written per second View Report: /report/oracle/PhysBlockIO.rpt

PhysicalReads	The number of physical reads. The value reflects requests to the operating system, which may not directly correspond to a physical read from a disk. [Sequential = SUM Non-Sequential = SUM]
PhysicalReads/s	The number of physical reads per second View Report: /report/oracle/PhysIO.rpt
PhysicalReadsDirect	The number of reads directly from disk, bypassing the buffer cache. For example, in high bandwidth, data-intensive operations such as parallel query, reads of disk blocks bypass the buffer cache to maximize transfer rates and to prevent the premature aging of shared data blocks resident in the buffer cache. [Sequential = SUM Non-Sequential = SUM]
PhysicalReadsDirect(LOB)	The number of large object writes directly from disk, bypassing the buffer cache [Sequential = SUM Non-Sequential = SUM]
PhysicalWrites	The number of physical writes. The value reflects requests to the operating system, which may not directly correspond to a physical write to a disk. [Sequential = SUM Non-Sequential = SUM]
PhysicalWrites/s	The number of physical writes per second View Report: /report/oracle/PhysIO.rpt
RecursiveCalls*	The number of recursive calls [Sequential = SUM Non-Sequential = SUM]
RecursiveToUserCallRatio	The ratio of recursive calls to user calls View Report: /report/oracle/CallRate.rpt
RedoEntries*	The number of times redo entries were copied into the redo log buffer [Sequential = SUM Non-Sequential = SUM]
RedoLogSpaceRequests*	The number of times a server process had to wait to acquire an entry in the redo log buffer [Sequential = SUM Non-Sequential = SUM]
RedoLogSpaceWaitRatio	The measure of redo log buffer memory allocation
Resource	The name of the Oracle Agent Instance [Sequential = ID Non-Sequential = ID]
RowSourceRatio	The percentage of total rows retrieved from full table scans View Report: /report/oracle/RowRatio.rpt
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
SortOverFlowRatio	The number of sorts that are using temporary segments View Report: /report/oracle/SortOverflow.rpt

System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
TableFetchByRowid*	The number of rows fetched using a ROWID [Sequential = SUM Non-Sequential = SUM]
TableFetchContinuedRow*	The number of rows fetched that span more than one database block [Sequential = SUM Non-Sequential = SUM]
TableScanRowsGotten*	The number of rows processed during a scan operation [Sequential = SUM Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
TransactionRate	The measure of application work View Report: /report/oracle/TransRate.rpt
UserCallRate	The measure of work being posed by client applications View Report: /report/oracle/CallRate.rpt
UserCalls*	The number of user calls (parse, execute, fetch) [Sequential = SUM Non-Sequential = SUM]
UserCallsPerParse	The measure of how well the application is managing context areas View Report: /report/oracle/CallRate.rpt
UserCommits*	The number of database transactions that were committed [Sequential = SUM Non-Sequential = SUM]
UserRollBackRatio	The rate at which application transactions are failing View Report: /report/oracle/Rollback.rpt
UserRollbacks*	The number of database transactions that were rolled back [Sequential = SUM Non-Sequential = SUM]

Class: Oracle
Subclass: Rollback Summary
IT Resource Name: /TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name: Oracle.Rollback Summary
Open Table Name: ORAROLLSUM
Resource: instance1, instance2, ...

Statistic Name:

RollGets	The number of rollback header gets [Sequential = SUM Non-Sequential = SUM]
RollWaits	The number of rollback header waits [Sequential = SUM Non-Sequential = SUM]
RollWrites	The number of bytes written to rollback segments [Sequential = SUM Non-Sequential = SUM]
RsSize	The size in bytes of the rollback segments [Sequential = LST Non-Sequential = SUM]

Class: Oracle
Subclass: Row
IT Resource Name: /TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name: Oracle.Row
Open Table Name: ORAROW
Resource: instance1, instance2, ...

Statistic Name:

RowCount	The total number of entries in the dictionary cache [Sequential = SUM Non-Sequential = SUM]
RowGets	The number of data requests for objects in the cache dictionary [Sequential = SUM Non-Sequential = SUM]
RowGetMisses	The number of data requests for objects in the dictionary cache that result in cache misses [Sequential = SUM Non-Sequential = SUM]
RowScans	The number of scan requests for objects in the cache [Sequential = SUM Non-Sequential = SUM]
RowScanMisses	The number of scan requests for objects in the dictionary cache that resulted in a cache miss [Sequential = SUM Non-Sequential = SUM]

Class:	Oracle
Subclass:	Session Summary
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Session Summary
Open Table Name:	ORASESSIONSUM
Resource:	instance1, instance2, ...
Statistic Name:	
ConnectTime	The connect time for the session in seconds [Sequential = SUM Non-Sequential = SUM]
CursorCacheCount	The total number of cursors cached. This statistic is incremented only if SESSION_CACHED_CURSORS > 0. [Sequential = LST Non-Sequential = SUM]
CursorCacheHits	The number of hits in the session cursor cache. A hit means that the SQL statement did not have to be reparsed. Subtract this statistic from parse count (total) to determine the real number of parses that occurred. [Sequential = SUM Non-Sequential = SUM]
LogicalReads	The sum of db block gets plus consistent gets [Sequential = SUM Non-Sequential = SUM]
PgaMemory	The current program global area (PGA) size for the session [Sequential = LST Non-Sequential = SUM]
PgaMemoryMax	The peak program global area (PGA) size for the session [Sequential = MAX Non-Sequential = MAX]
StoredProcedure Space	The amount of memory the session is using for stored procedures [Sequential = LST Non-Sequential = SUM]
UgaMemory	The current user global area (UGA) size for the session [Sequential = LST Non-Sequential = SUM]
UgaMemoryMax	The peak user global area (UGA) size for the session [Sequential = MAX Non-Sequential = MAX]
Class:	Oracle
Subclass:	System
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.System
Open Table Name:	ORASYS
Resource:	instance1, instance2, ...
Statistic Name:	
AvgLogons	The average number of logons during the interval [Sequential = AVG Non-Sequential = SUM]
AvgProcesses	The average number of processes during the interval [Sequential = AVG Non-Sequential = SUM]

AvgSessions	The average number of sessions during the interval [Sequential = AVG Non-Sequential = SUM]
BgProcesses	The number of background processes [Sequential = LST Non-Sequential = SUM]
CachedCommitSCN	The number of times the system change number (SCN) of a commit operation was cached for fast retrieval by other Oracle tasks. This statistic is useful only for internal debugging purposes. [Sequential = SUM Non-Sequential = SUM]
CallsToGetSnapshotKCMGSS	The number of times a snapshot system change number (SCN) was allocated. The SCN is allocated at the start of a transaction. [Sequential = SUM Non-Sequential = SUM]
CallsToKCMGAS	The number of calls to routine kcmgas to get a new system change number (SCN) [Sequential = SUM Non-Sequential = SUM]
CallsToKCMGCS	The number of calls to routine kcmgcs to get a current system change number (SCN) [Sequential = SUM Non-Sequential = SUM]
CallsToKCMGRS	The number of calls to routine kcmgrs to get a recent system change number (SCN) [Sequential = SUM Non-Sequential = SUM]
ChangeWriteTime	The elapsed redo write time for changes made to CURRENT blocks in seconds. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
CleanoutsAndRollbacks	The number of consistent gets that require both block rollbacks and block cleanouts [Sequential = SUM Non-Sequential = SUM]
CleanoutsOnly	The number of consistent gets that require only block cleanouts, no rollbacks [Sequential = SUM Non-Sequential = SUM]
ClusterKeyScanGets	The number of blocks obtained in a cluster scan [Sequential = SUM Non-Sequential = SUM]
ClusterKeyScans	The number of cluster scans that were started [Sequential = SUM Non-Sequential = SUM]
ColdRecycleReads	The number of buffers that were read through the least recently used end of the recycle cache with fast aging strategy [Sequential = SUM Non-Sequential = SUM]
CommitClnoutFailuresBlkLost	The number of times Oracle attempted a cleanout at commit but could not find the correct block due to forced write, replacement, or switch CURRENT [Sequential = SUM Non-Sequential = SUM]
CommitClnoutFailuresBuffer	The number of times Oracle attempted a cleanout at commit, but the buffer was currently being written [Sequential = SUM Non-Sequential = SUM]

CommitClnoutFailuresCallBack	The number of times the cleanout callback function returns FALSE [Sequential = SUM Non-Sequential = SUM]
CommitClnoutFailuresCannotpin	The total number of times a commit cleanout was performed but failed because the block could not be pinned [Sequential = SUM Non-Sequential = SUM]
CommitClnoutFailuresHotbkup	The number of times Oracle attempted block cleanout at commit during hot backup. The image of the block needs to be logged before the buffer can be made dirty. [Sequential = SUM Non-Sequential = SUM]
CommitClnoutFailuresWrite	The number of times a cleanout block at commit was performed but the writes to the database had been temporarily disabled [Sequential = SUM Non-Sequential = SUM]
CommitClnouts	The total number of times the cleanout block of a commit function was performed [Sequential = SUM Non-Sequential = SUM]
CommitClnoutsCompleted	The number of times the cleanout block of a commit function completed successfully [Sequential = SUM Non-Sequential = SUM]
CommitSCNcached	The number of times the system change number of a commit function was cached [Sequential = SUM Non-Sequential = SUM]
ConsistentChanges	The number of consistent get operations that cannot accept the block in its current state [Sequential = SUM Non-Sequential = SUM]
ConsistentGets	The number of requests to the buffer manager to locate a database block as part of a consistent-read operation [Sequential = SUM Non-Sequential = SUM]
CpuUsed	The amount of CPU time in centiseconds (1/100th of a second) used by the Oracle instance. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM] View Report: /report/oracle/CpuUsed.rpt
CpuUsed (seconds)	The amount of CPU time in seconds used by the Oracle instance. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
CrBlocksCreated	The number of CURRENT blocks cloned to create consistent-read (CR) blocks. The most common reason for cloning is that the buffer is held in a incompatible mode. [Sequential = SUM Non-Sequential = SUM]
CurrentBlksConvForCR	The number of CURRENT blocks converted to CR state [Sequential = SUM Non-Sequential = SUM]

CursorAuthentications	The number of privilege checks conducted during execution of an operation [Sequential = SUM Non-Sequential = SUM]
DataBlksConsistentReads	The number of undo records applied to data blocks that have been rolled back for consistent-read (CR) purposes [Sequential = SUM Non-Sequential = SUM]
DbBlockChanges	The number of database blocks in memory that created a “dirty block” [Sequential = SUM Non-Sequential = SUM]
DbBlockGets	The number of requests to the buffer manager for a database block regardless of read consistency [Sequential = SUM Non-Sequential = SUM]
DDLStmntsParallel	The number of data definition language (DDL) statements that were executed in parallel [Sequential = SUM Non-Sequential = SUM]
DeferredBlkCleanout	The number of times cleanout records are deferred within the data buffer [Sequential = SUM Non-Sequential = SUM]
DFOTreesParallel	The number of times a serial execution plan was converted to a parallel plan [Sequential = SUM Non-Sequential = SUM]
DirtyBuffersInspected	The number of dirty buffers found by the user process while it is looking for a buffer to reuse [Sequential = SUM Non-Sequential = SUM]
DiskSorts	The number of sorts that allocated work space on disk [Sequential = SUM Non-Sequential = SUM]
DMLStmntsParallel	The number of data manipulation language (DML) statements that were executed in parallel [Sequential = SUM Non-Sequential = SUM]
ExchangeDeadlocks	The number of times a process detected a potential deadlock when exchanging two buffers and raised an internal, restartable error. Index scans are the only operations that perform exchanges. [Sequential = SUM Non-Sequential = SUM]
ExecuteCount	The total number of calls (user and recursive) that executed SQL statements [Sequential = SUM Non-Sequential = SUM]
FreeBufferInspected	The number of buffers skipped over from the end of a least recently used (LRU) queue to find a reusable buffer. The difference between this statistic and the DirtyBuffersInspected statistic that could not be used because the buffer had a user, a waiter, or were being read or written, or because the buffers were busy or needed to be written after rapid aging out. [Sequential = SUM Non-Sequential = SUM]

FreeBufferRequested	The number of times a reusable buffer or a free buffer was requested to create or load a block [Sequential = SUM Non-Sequential = SUM]
HotBuffersToHeadLRU	The number of times Oracle moves a hot buffer to the head of its replacement list to keep it from being reused once the buffer has reached the tail of its replacement list [Sequential = SUM Non-Sequential = SUM]
ImmediateCRBlkClnout	The number of times cleanout records are applied immediately during consistent-read requests [Sequential = SUM Non-Sequential = SUM]
ImmediateCurrBlkClnout	The number of times cleanout records are applied immediately during current gets. Compare this statistic with deferred (CURRENT) block cleanout applications. [Sequential = SUM Non-Sequential = SUM]
IndexFastFullscanDirRead	The number of fast full scans initiated using direct read [Sequential = SUM Non-Sequential = SUM]
IndexFastFullscanFull	The number of fast full scans initiated for full segments [Sequential = SUM Non-Sequential = SUM]
IndexFastFullscanRowid	The number of fast full scans initiated with row-id endpoints specified [Sequential = SUM Non-Sequential = SUM]
InstanceRecoveryFreezeCnt	The number of times the database is frozen during instance recovery. This statistic is not available with Oracle 11.2 and later. [Sequential = SUM Non-Sequential = SUM]
KCMCCSCalledGet	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
KCMGSSWaitedForBatching	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
LeafNodeSplits	The number of times an index leaf node was split because of the insertion of an additional value [Sequential = SUM Non-Sequential = SUM]
Logons	The number of current logons [Sequential = LST Non-Sequential = SUM]
LogonsCumulative	The total number of logons since the instance started. This statistic is useful only in V\$SYSSTAT. It gives an instance overview of all processes that logged on. [Sequential = SUM Non-Sequential = SUM]
MaxLogons	The maximum number of logons during the interval [Sequential = MAX Non-Sequential = SUM]
MaxProcesses	The maximum number of processes during the interval [Sequential = MAX Non-Sequential = SUM]
MaxSessions	The maximum number of sessions during the interval [Sequential = MAX Non-Sequential = SUM]

MemorySorts	The number of sorts that did not require allocation of work space on disk [Sequential = SUM Non-Sequential = SUM]
MessagesReceived	The number of messages received between background processes [Sequential = SUM Non-Sequential = SUM]
MessagesSent	The number of messages sent between background processes [Sequential = SUM Non-Sequential = SUM]
NativeHashExecute	The number of hash operations performed using native arithmetic rather than Oracle NUMBERs [Sequential = SUM Non-Sequential = SUM]
NativeHashFail	The number of hash operations performed using native arithmetic that failed, requiring the hash operation to be performed with Oracle NUMBERs [Sequential = SUM Non-Sequential = SUM]
NoBufferToKeepPinCnt	The number of times a visit to a buffer was attempted, but the buffer was not found where expected. This statistic is useful only for internal debugging purposes. [Sequential = SUM Non-Sequential = SUM]
NoWorkConsistentReadgets	The number of consistent gets that require neither block cleanouts nor rollbacks [Sequential = SUM Non-Sequential = SUM]
Objects	The number of database objects [Sequential = LST Non-Sequential = SUM]
OpenCursors	The number of open cursors [Sequential = LST Non-Sequential = SUM] View Report: /report/oracle/Cursor.rpt
OpenCursorsCumulative	The total number of cursors opened since the instance started [Sequential = LST Non-Sequential = SUM]
OpensCacheReplacement	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
OpensOfReplacedFiles	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
ParseCount	The number of parse calls received from connected applications and recursive calls [Sequential = SUM Non-Sequential = SUM]
ParseCountFailures	The number of parses that failed to parse [Sequential = SUM Non-Sequential = SUM]
ParseCountHard	The total number of parse calls (real parses). A hard parse is a very expensive operation in terms of memory use, because it requires Oracle to allocate a workheap and other memory structures and then build a parse tree. [Sequential = SUM Non-Sequential = SUM]

ParseTimeCpu	The total CPU time used for parsing (hard and soft) in seconds. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
ParseTimeElapsed	The total elapsed time for parsing in seconds. Subtract parse time cpu from this statistic to determine the total waiting time for parse resources. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
PctProcessCapacity*	The percentage of process capacity currently used View Report: /report/oracle/PctConnections.rpt
PctSessionsCapacity*	The percentage of session capacity currently used View Report: /report/oracle/PctConnections.rpt
PrefetchedBlocks	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
PrefetchedBlocksAgedout	The number of contiguous and noncontiguous blocks that were prefetched but aged out before use [Sequential = SUM Non-Sequential = SUM]
Processes	The number of current active processes [Sequential = LST Non-Sequential = SUM] View Report: /report/oracle/Processes.rpt
ProcessesLimit	The maximum number of operating system user processes that can simultaneously connect to the Oracle instance [Sequential = LST Non-Sequential = SUM] View Report: /report/oracle/Processes.rpt
ProcessLastNonidleTime	The last time this process was executed. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = LST Non-Sequential = SUM]
PXlocalmsgrecv	The number of local messages received for parallel execution within the instance local to the current session [Sequential = SUM Non-Sequential = SUM]
PXlocalmsgsent	The number of local messages sent for parallel execution within the instance local to the current session [Sequential = SUM Non-Sequential = SUM]
PXremotemsgrecv	The number of remote messages received for parallel execution within the instance local to the current session [Sequential = SUM Non-Sequential = SUM]
PXRemotemsgsent	The number of remote messages sent for parallel execution within the instance local to the current session [Sequential = SUM Non-Sequential = SUM]

QueriesParallelized	The number of SELECT statements executed in parallel [Sequential = SUM Non-Sequential = SUM]
RecoveryArrayReads	The number of reads performed during recovery [Sequential = SUM Non-Sequential = SUM]
RecoveryArrayReadTime	The elapsed time of I/O in seconds during recovery [Sequential = SUM Non-Sequential = SUM]
RecoveryBlocksRead	The number of blocks read during recovery [Sequential = SUM Non-Sequential = SUM]
RecursiveCalls	The number of recursive calls [Sequential = SUM Non-Sequential = SUM]
RecursiveCpuUsage	The total CPU time in seconds used by non-user calls (recursive calls). Subtract this value from CPU used by the session to determine how much CPU time was used by the user calls. [Sequential = SUM Non-Sequential = SUM]
RedoBlocksWritten	The total number of redo blocks written. This statistic divided by redo writes equals number of blocks per write. [Sequential = SUM Non-Sequential = SUM]
RedoBufferAllocationRetries	The total number of retries necessary to allocate space in the redo buffer. Retries are needed either because the redo writer has fallen behind or because an event such as a log switch is occurring. [Sequential = SUM Non-Sequential = SUM]
RedoEntries	The number of times redo entries were copied into the redo log buffer [Sequential = SUM Non-Sequential = SUM]
RedoLogSpaceRequests	The number of times a server process had to wait to acquire an entry in the redo log buffer [Sequential = SUM Non-Sequential = SUM]
RedoLogSpaceWaitTime	The total elapsed waiting time for redo log space requests in seconds [Sequential = SUM Non-Sequential = SUM]
RedoLogSwitchInterrupts	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
RedoOrderingMarks	The number of times a system change number (SCN) was allocated to force a redo record to have a higher SCN than a record generated in another thread using the same block [Sequential = SUM Non-Sequential = SUM]
RedoSize	The total amount of redo generated in bytes [Sequential = SUM Non-Sequential = SUM]
RedoSyncTime	The elapsed time of all redo sync write calls in seconds. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]

RedoSyncWrites	<p>The number of times a change being applied to the log buffer must be written out to disk due to a commit. The log buffer is a circular buffer that the log writer (LGWR) periodically flushes. Usually, redo that is generated and copied into the log buffer does not need to be flushed out to disk immediately.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RedoWastage	<p>The number of bytes wasted because redo blocks needed to be written before they are completely full. Early writing may be needed to commit transactions, to be able to write a database buffer, or to switch logs.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RedoWriterLatchingTime	<p>The elapsed time in seconds needed by the log writer (LGWR) to obtain and release each copy latch. This statistic is not available with Oracle 11.2 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RedoWrites	<p>The total number of writes by the log writer (LGWR) to the redo log files. Redo blocks written divided by this statistic equals the number of blocks per write.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RedoWriteTime	<p>The total elapsed time of the write from the redo log buffer to the current redo log file in seconds. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RemoteInstanceUndoBlkWrites	<p>This statistic is not available for the Oracle Database Server. The value is reported as <N/A>.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RemoteInstanceUndoHdrWrites	<p>This statistic is not available for the Oracle Database Server. The value is reported as <N/A>.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
Resource	<p>The name of the Oracle Agent Instance</p> <p>[Sequential = ID Non-Sequential = ID]</p>
RollbackChangesUndo	<p>The number of undo records applied to user-requested rollback changes (not consistent-read rollbacks)</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RollbacksOnly	<p>The number of consistent gets that require only block rollbacks, no block cleanouts</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RowsFetchedViaCallback	<p>The rows fetched using callback. This statistic is useful primarily for internal debugging purposes.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
RowSorts	<p>The total number of rows sorted</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
SerializableAborts	<p>The number of times an SQL statement in a serializable isolation level had to abort</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>

Sessions	The number of current sessions [Sequential = LST Non-Sequential = SUM] View Report: /report/oracle/Sessions.rpt
SessionsLimit	The maximum number of user and system sessions allowed for the Oracle instance [Sequential = LST Non-Sequential = SUM] View Report: /report/oracle/Sessions.rpt
SummedDirtyQueueLength	The sum of the dirty least recently used (LRU) queue length after every write request. Divide by write requests to get the average queue length after write completion. [Sequential = SUM Non-Sequential = SUM]
SwitchCurrToNewBuffer	The number of times the CURRENT block moved to a different buffer, leaving a consistent-read (CR) block in the original buffer [Sequential = SUM Non-Sequential = SUM]
TableFetchByRowid	The number of rows fetched using a ROWID [Sequential = SUM Non-Sequential = SUM]
TableFetchContinuedRow	The number of rows fetched that span more than one database block [Sequential = SUM Non-Sequential = SUM]
TableScanBlocksGotten	The number of blocks encountered during a scan operation, in which each row is retrieved sequentially by Oracle [Sequential = SUM Non-Sequential = SUM]
TableScanRowsGotten	The number of rows processed during a scan operation [Sequential = SUM Non-Sequential = SUM]
TableScansCachePart	The number of range scans performed on tables that have the CACHE option enabled [Sequential = SUM Non-Sequential = SUM]
TableScansDirectRead	The number of table scans performed with direct read (bypassing the buffer cache) [Sequential = SUM Non-Sequential = SUM]
TableScansLong	The number of table scans for long tables [Sequential = SUM Non-Sequential = SUM]
TableScansRowidRanges	During parallel query, the number of table scans conducted with specified ROWID ranges [Sequential = SUM Non-Sequential = SUM]
TableScansShort	The number of table scans for short tables [Sequential = SUM Non-Sequential = SUM]
TotalFileOpens	This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
TransactionRollbacks	The number of transactions being successfully rolled back [Sequential = SUM Non-Sequential = SUM]

TransactionTblReadRollbacks	The number of times rollback segment headers are rolled back to create consistent-read (CR) blocks [Sequential = SUM Non-Sequential = SUM]
TransactionTblReadUndo	The number of undo records applied to transaction tables that have been rolled back for consistent-read (CR) purposes [Sequential = SUM Non-Sequential = SUM]
UserCalls	The number of user calls (parse, execute, fetch) [Sequential = SUM Non-Sequential = SUM]
UserCommits	The number of database transactions that were committed [Sequential = SUM Non-Sequential = SUM]
UserRollbacks	The number of database transactions that were rolled back [Sequential = SUM Non-Sequential = SUM]
WriteClonesBackground	The number of times a background process clones a CURRENT buffer that is being written. The clone becomes the new, accessible CURRENT buffer, leaving the original buffer (now the clone) to complete writing. [Sequential = SUM Non-Sequential = SUM]
WriteClonesForeground	The number of times a foreground process clones a CURRENT buffer that is being written. The clone becomes the new, accessible CURRENT buffer, leaving the original buffer (now the clone) to complete writing. [Sequential = SUM Non-Sequential = SUM]

Class:	Oracle
Subclass:	Wait
IT Resource Name:	<i>/TeamQuest/System/systemname/Oracle/instancename</i>
TeamQuest Table Name:	Oracle.Wait
Open Table Name:	ORAWAIT
Resource:	instance1, instance2, ...
Statistic Name:	
SystemUndoBlock	The number of waits for buffers containing blocks from the SYSTEM rollback segment [Sequential = SUM Non-Sequential = SUM]
SystemUndoHeader	The number of waits for buffers containing header blocks from the SYSTEM rollback segment [Sequential = SUM Non-Sequential = SUM]
UndoBlock	The number of waits for buffers containing blocks other than those from the SYSTEM rollback segment [Sequential = SUM Non-Sequential = SUM]
UndoHeader	The number of waits for buffers containing header blocks other than those from the SYSTEM rollback segment [Sequential = SUM Non-Sequential = SUM]
Class:	Oracle
Subclass:	N/A
IT Resource Name:	<i>/TeamQuest/System/systemname/Oracle/instancename</i>
TeamQuest Table Name:	Oracle
Open Table Name:	ORA
Resource:	instance1, instance2, ...
Statistic Name:	
Etime	The seconds elapsed between two data samples of the Oracle Data Agent [Sequential = SUM Non-Sequential = SUM]
tqorap_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqorap_interval	The seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

12.7. Latch Statistics

The Oracle.Latch table stores detailed information about latches within the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	Latch
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Oracle/ <i>instancename</i>
TeamQuest Table Name:	Oracle.Latch
Open Table Name:	ORALATCH
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
gets	The number of times the latch was obtained with a wait during the interval [Sequential = SUM Non-Sequential = SUM]
gets_t	The total number of times the latch was obtained with a wait since the instance started [Sequential = LST Non-Sequential = SUM]
imm_gets	The number of times the latch was obtained without a wait during the interval [Sequential = SUM Non-Sequential = SUM]
imm_gets_t	The total number of times the latch was obtained without a wait since the instance started [Sequential = LST Non-Sequential = SUM]
imm_miss_ratio	The ratio of imm_misses to imm_gets during the interval [Sequential = AVG Non-Sequential = AVG]
imm_misses	The number of times the latch failed during the interval [Sequential = SUM Non-Sequential = SUM]
imm_misses_t	The total number of times the latch failed to be obtained without a wait since the instance started [Sequential = LST Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]

level_num	The latch level number [Sequential = ID Non-Sequential = ID]
miss_ratio	The ratio of misses to gets during the interval [Sequential = AVG Non-Sequential = AVG]
misses	The number of times the latch was obtained with a wait but failed on the first try during the interval [Sequential = SUM Non-Sequential = SUM]
misses_t	The total number of times the latch was obtained with a wait but failed on the first try since the instance started [Sequential = LST Non-Sequential = SUM]
name	The name of the latch [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
sleeps	The number of sleeps that occurred when trying to obtain the latch with a wait during the interval [Sequential = SUM Non-Sequential = SUM]
sleeps_t	The total number of sleeps that occurred when trying to obtain the latch with a wait since the instance started [Sequential = LST Non-Sequential = SUM]
spin_gets	The number of gets that missed on the first try but succeeded on spin during the interval [Sequential = SUM Non-Sequential = SUM]
spin_gets_t	The total number of gets that missed on the first try but succeeded on spin since the instance started [Sequential = LST Non-Sequential = SUM]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
waiters_woken	The number of times a wait for the latch was awakened during the interval [Sequential = SUM Non-Sequential = SUM]
waiters_woken_t	The total number of times a wait for the latch was awakened since the instance started [Sequential = LST Non-Sequential = SUM]

waits_holding_latch	The number of waits for the latch that held a different latch during the interval [Sequential = SUM Non-Sequential = SUM]
waits_holding_latch_t	The total number of waits for the latch that held a different latch since the instance started [Sequential = LST Non-Sequential = SUM]

12.8. Library Cache Statistics

The Oracle.LibraryCache table stores detailed information about the library cache of the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	LibraryCache
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.LibraryCache
Open Table Name:	ORALIBCACHE
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
get_hit_ratio	The ratio of get hits to gets during the interval [Sequential = AVG Non-Sequential = AVG]
get_hits	The number of times a handle of an object for the namespace was found in memory during the interval [Sequential = SUM Non-Sequential = SUM]
get_hits_t	The total number of times a handle of an object for the namespace was found in memory since the instance started [Sequential = LST Non-Sequential = SUM]
gets	The number of times a lock request was performed by objects of the namespace during the interval [Sequential = SUM Non-Sequential = SUM]
gets_t	The total number of times a lock request was performed by objects of the namespace since the instance started [Sequential = LST Non-Sequential = SUM]

Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
invalidations	The number of times an object of the namespace was marked invalid during the interval [Sequential = SUM Non-Sequential = SUM]
invalidations_t	The total number of times an object of the namespace was marked invalid since the instance started [Sequential = LST Non-Sequential = SUM]
namespace	The library cache namespace [Sequential = ID Non-Sequential = ID]
pin_hit_ratio	The ratio of pin hits during the interval [Sequential = AVG Non-Sequential = AVG]
pin_hits	The number of times all of the metadata pieces for the given object of the namespace were found in memory during the interval [Sequential = SUM Non-Sequential = SUM]
pin_hits_t	The total number of times all of the metadata pieces for the given object of the namespace were found in memory since the instance started [Sequential = LST Non-Sequential = SUM]
pins	The number of times a pin was requested by objects of the namespace during the interval [Sequential = SUM Non-Sequential = SUM]
pins_t	The total number of times a pin was requested by objects of the namespace since the instance started [Sequential = LST Non-Sequential = SUM]
reloads	The number of reloads performed during a pin request by objects of the namespace during the interval [Sequential = SUM Non-Sequential = SUM]
reloads_t	The total number of reloads performed during a pin request by objects of the namespace since the instance started [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

12.9. Listener Alarm Statistics

The Oracle.Alarm_Listener table stores alarm information about the status of Oracle listeners.

Table Field Hierarchy

Class:	Oracle
Subclass:	Alarm_Listener
IT Resource Name:	/TeamQuest/System/systemname/Oracle
TeamQuest Table Name:	Oracle.Alarm_Listener
Open Table Name:	ORAALARMLISTENER
Collection interval:	60 seconds (default)
Default retention:	30 days
Table type:	Performance

Statistic Name	Description
Alarm_ID	The user-defined alarm identifier assigned to the alarm [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Listener	The name of the listener from which the data is obtained. Up to 23 characters are displayed. [Sequential = ID Non-Sequential = ID]
Match_ID	A tab-delimited concatenation of the time consolidation fields for the record. This is a hidden field and is for internal use only. [Sequential = ID Non-Sequential = ID]
Severity	The severity of the alarm (Normal, Warning, Minor, Major, or Critical) [Sequential = LST Non-Sequential = ID]
status	The status of the instance (Up or Down) [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

12.10. Lock Alarm Statistics

The Oracle.Alarm_Lock table stores alarm information about locks within an Oracle instance.

Table Field Hierarchy

Class:	Oracle
Subclass:	Alarm_Lock
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Alarm_Lock
Open Table Name:	ORAALARMLOCK
Collection interval:	60 seconds (default)
Default retention:	30 days
Table type:	Performance

Statistic Name	Description
Alarm_ID	The user-defined alarm identifier assigned to the alarm [Sequential = ID Non-Sequential = ID]
block_status	Specifies if the lock is blocking another lock (yes or no) [Sequential = LST Non-Sequential = ID]
curr_lock_mode	The current lock mode [Sequential = LST Non-Sequential = ID]
id1	The first part of the type-specific lock identifier [Sequential = ID Non-Sequential = ID]
id2	The second part of the type-specific lock identifier [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
lock_time	The amount of time in seconds that the lock has been in its current mode [Sequential = LST Non-Sequential = SUM]
Match_ID	A tab-delimited concatenation of the time consolidation fields for the record. This is a hidden field and is for internal use only. [Sequential = ID Non-Sequential = ID]
owner	The database user that owns the lock [Sequential = ID Non-Sequential = ID]
req_lock_mode	The requested lock mode [Sequential = LST Non-Sequential = ID]
serial_num	The session serial number for the owner of the lock [Sequential = ID Non-Sequential = ID]

Severity	The severity of the alarm (Normal, Warning, Minor, Major, or Critical) [Sequential = LST Non-Sequential = ID]
sid	The session identifier for the owner of the lock [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
type	The type of lock [Sequential = ID Non-Sequential = ID]

12.11. Rollback Segment Statistics

The Oracle.Rollback table stores detailed information about the rollback segments used by the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	Rollback
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Rollback
Open Table Name:	ORAROLL
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
extends	The number of times the rollback segment extended during the interval [Sequential = SUM Non-Sequential = SUM]
extends_t	The total number of times the rollback segment extended since the instance started [Sequential = LST Non-Sequential = SUM]
extents	The current number of extents in the rollback segment [Sequential = LST Non-Sequential = SUM]

gets	The number of header gets for the rollback segment during the interval [Sequential = SUM Non-Sequential = SUM]
gets_t	The total number of header gets for the rollback segment since the instance started [Sequential = LST Non-Sequential = SUM]
hwmsize	The maximum size in megabytes of the rollback segment since the instance started [Sequential = MAX Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
name	The name of the rollbacks segment [Sequential = ID Non-Sequential = ID]
optsize	The optimal size in megabytes of the rollback segment [Sequential = LST Non-Sequential = SUM]
rssize	The current size in megabytes of the rollback segment [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shrinks	The number of times the size of the rollback segment decreased during the interval [Sequential = SUM Non-Sequential = SUM]
shrinks_t	The total number of times the size of the rollback segment decreased since the instance started [Sequential = LST Non-Sequential = SUM]
status	The status of the rollback segment [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
usn	The rollback segment number [Sequential = ID Non-Sequential = ID]
waits	The number of header waits for the rollback segment during the interval [Sequential = SUM Non-Sequential = SUM]
waits_t	The total number of header waits for the rollback segment since the instance started [Sequential = LST Non-Sequential = SUM]

wraps	The number of times the rollback segment wrapped during the interval [Sequential = SUM Non-Sequential = SUM]
wraps_t	The total number of times the rollback segment wrapped since the instance started [Sequential = LST Non-Sequential = SUM]
writes	The number of kilobytes written to the rollback segment during the interval [Sequential = SUM Non-Sequential = SUM]
writes_t	The total number of kilobytes written to the rollback segment since the instance started [Sequential = LST Non-Sequential = SUM]
xacts	The current number of active transactions for the rollback segment [Sequential = LST Non-Sequential = SUM]

12.12. Row Cache Statistics

The Oracle.RowCache table stores detailed information about the row cache of the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	RowCache
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Oracle/ <i>instancename</i>
TeamQuest Table Name:	Oracle.RowCache
Open Table Name:	ORAROWCACHE
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
cache_num	The row cache identifier number [Sequential = ID Non-Sequential = ID]
count	The total number of entries in the cache for the interval [Sequential = LST Non-Sequential = SUM]
fixed	The number of fixed entries in the cache for the interval [Sequential = LST Non-Sequential = SUM]
flushes	The number of flushes to disk during the interval [Sequential = SUM Non-Sequential = SUM]

flushes_t	The total number of flushes to disk since the instance started [Sequential = LST Non-Sequential = SUM]
get_miss_ratio	The ratio of get_misses to gets during the interval [Sequential = AVG Non-Sequential = AVG]
get_misses	The number of requests for objects in the cache that resulted in cache misses during the interval [Sequential = SUM Non-Sequential = SUM]
get_misses_t	The total number of requests for objects in the cache that resulted in cache misses since the instance started [Sequential = LST Non-Sequential = SUM]
gets	The number of data requests for objects in the cache during the interval [Sequential = SUM Non-Sequential = SUM]
gets_t	The total number of data requests for objects in the cache since the instance started [Sequential = LST Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
modifications	The number of inserts, updates, or deletions for objects in the cache during the interval [Sequential = SUM Non-Sequential = SUM]
modifications_t	The total number of inserts, updates, or deletions for objects in the cache since the instance started [Sequential = LST Non-Sequential = SUM]
parameter	The name of the initialization parameter that determines the number of entries in the data dictionary cache [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
scan_miss_ratio	The ratio of scan_misses to scans during the interval [Sequential = AVG Non-Sequential = AVG]
scan_misses	The number of scan requests for objects in the cache that resulted in cache misses during the interval [Sequential = SUM Non-Sequential = SUM]
scan_misses_t	The total number of scan requests for objects in the cache that resulted in cache misses since the instance started [Sequential = LST Non-Sequential = SUM]
scans	The number of scan requests for objects in the cache during the interval [Sequential = SUM Non-Sequential = SUM]

scans_t	The total number of scan requests for objects in the cache since the instance started [Sequential = LST Non-Sequential = SUM]
sub_num	The subordinate set number [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
type	The parent or subordinate row cache type [Sequential = ID Non-Sequential = ID]
usage	The number of cache entries that contain valid data for the interval [Sequential = LST Non-Sequential = SUM]

12.13. Segment Alarm Statistics

The Oracle.Alarm_Segment table stores alarm information about segments within an Oracle database.

Table Field Hierarchy

Class:	Oracle
Subclass:	Alarm_Segment
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Oracle/ <i>instancename</i>
TeamQuest Table Name:	Oracle.Alarm_Segment
Open Table Name:	ORAALARMSEGMENT
Collection interval:	60 seconds (default)
Default retention:	30 days
Table type:	Performance

Statistic Name	Description
Alarm_ID	The user-defined alarm identifier assigned to the alarm [Sequential = ID Non-Sequential = ID]
can_extend	Specifies if there is enough space for another extent to be allocated for the segment within the tablespace (yes or no) [Sequential = LST Non-Sequential = ID]
extents	The number of extents currently used by the segment [Sequential = LST Non-Sequential = SUM]
extents_avail	The number of extents below the maximum that may be allocated for the segment. This is the value of the max_extents field minus the value of the extents field. [Sequential = LST Non-Sequential = SUM]

Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Match_ID	A tab-delimited concatenation of the time consolidation fields for the record. This is a hidden field and is for internal use only. [Sequential = ID Non-Sequential = ID]
max_extents	The maximum number of extents allowed for the segment [Sequential = LST Non-Sequential = SUM]
max_free_extent	The size in megabytes of the largest free extent within the tablespace that contains the segment [Sequential = LST Non-Sequential = MAX]
next_extent	The size in megabytes of the next extent to be allocated for the segment [Sequential = LST Non-Sequential = SUM]
owner	The database users that owns the segment [Sequential = ID Non-Sequential = ID]
segment_name	The name of the segment [Sequential = ID Non-Sequential = ID]
segment_type	The type of segment [Sequential = ID Non-Sequential = ID]
Severity	The severity of the alarm (Normal, Warning, Minor, Major, or Critical) [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
tablespace_name	The name of the tablespace in which the segment resides [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

12.14. Session Statistics

The Oracle.Session table stores detailed information about current sessions for the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	Session
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.Session
Open Table Name:	ORASESSION
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avg_waittime	The average amount of time in seconds spent waiting by the session during the interval. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = AVG Non-Sequential = AVG]
block_gets	The number of database block gets performed by the session during the interval. The sum of this value with the cons_gets value is a measure of logical reads. [Sequential = SUM Non-Sequential = SUM]
block_gets_t	The total number of database block gets performed by the session since it began. The sum of this value with the cons_gets_t value is a measure of logical reads. [Sequential = LST Non-Sequential = SUM]
clnt_pid	The operating system process identifier for the client program using the session [Sequential = ID Non-Sequential = ID]
clnt_program	The operating system name for the client program using the session [Sequential = ID Non-Sequential = ID]
clnt_system	The name of the system on which the client program using the session is running [Sequential = ID Non-Sequential = ID]
clnt_term	The operating system terminal name for the client program using the session [Sequential = ID Non-Sequential = ID]

clnt_user	The operating system user name for the client program using the session [Sequential = ID Non-Sequential = ID]
cons_gets	The number of consistent gets performed by the session during the interval. The sum of this value with the block_gets value is a measure of logical reads. [Sequential = SUM Non-Sequential = SUM]
cons_gets_t	The total number of consistent gets performed by the session since it began. The sum of this value with the block_gets_t value is a measure of logical reads. [Sequential = LST Non-Sequential = SUM]
cpu_used	The number of CPU seconds used by the session over the interval. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
cpu_used_t	The total number of CPU seconds used by the session since it began. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = LST Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
logon_time	The logon time of the session [Sequential = FST Non-Sequential = ID]
ora_pid	The Oracle process identifier for the Oracle process related to the session [Sequential = ID Non-Sequential = ID]
ora_user	The Oracle user name that is using the session [Sequential = ID Non-Sequential = ID]
parse_cnt	The number of parse calls made by the session during the interval [Sequential = SUM Non-Sequential = SUM]
parse_cnt_t	The total number of parse calls made by the session since it began [Sequential = LST Non-Sequential = SUM]
pga_mem	The current size in kilobytes of the program global area (PGA) for the session [Sequential = LST Non-Sequential = SUM]
phys_rds	The number of physical reads performed by the session during the interval. The value reflects requests to the operating system, which may not directly correspond to a physical read from a disk. [Sequential = SUM Non-Sequential = SUM]
phys_rds_t	The total number of physical reads performed by the session since it began. The value reflects requests to the operating system, which may not directly correspond to a physical read from a disk. [Sequential = LST Non-Sequential = SUM]

phys_wrts	The number of physical writes performed by the session during the interval. The value reflects requests to the operating system and may not directly correspond to a physical write to a disk. [Sequential = SUM Non-Sequential = SUM]
phys_wrts_t	The total number of physical writes performed by the session since it began. The value reflects requests to the operating system and may not directly correspond to a physical write to a disk. [Sequential = LST Non-Sequential = SUM]
rec_calls	The number of recursive calls made by the session during the interval. Recursive calls are a measure of internal work done by the Oracle instance while performing requests from a user. [Sequential = SUM Non-Sequential = SUM]
rec_calls_t	The total number of recursive calls made by the session since it began. Recursive calls are a measure of internal work done by the Oracle instance while performing requests from a user. [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
serial_num	The session serial number [Sequential = ID Non-Sequential = ID]
server	The server type for the session [Sequential = ID Non-Sequential = ID]
sid	The session identifier [Sequential = ID Non-Sequential = ID]
sql_addr	The address of the SQL cursor currently being used by the session [Sequential = ID Non-Sequential = ID]
sql_hash_value	The value that is used with sql_addr to uniquely identify the SQL cursor currently being used by the session [Sequential = ID Non-Sequential = ID]
sql_text	The first 64 characters of text for the SQL cursor currently being used by the session [Sequential = ID Non-Sequential = ID]
svr_pid	The operating system process identifier for the Oracle process related to the session [Sequential = ID Non-Sequential = ID]
svr_program	The operating system name for the Oracle process related to the session [Sequential = ID Non-Sequential = ID]
svr_term	The operating system terminal name for the Oracle process related to the session [Sequential = ID Non-Sequential = ID]
svr_user	The operating system user name for the Oracle process related to the session [Sequential = ID Non-Sequential = ID]

status	The status of the session [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
timeouts	The number of timeouts by the session during the interval [Sequential = SUM Non-Sequential = SUM]
timeouts_t	The total number of timeouts by the session since it began. [Sequential = LST Non-Sequential = SUM]
type	The session type [Sequential = ID Non-Sequential = ID]
uga_mem	The current size in kilobytes of the user global area (UGA) for the session [Sequential = LST Non-Sequential = SUM]
usr_calls	The number of user calls made by the session during the interval. User calls are incremented whenever a parse, execute, or fetch is performed by the user. [Sequential = SUM Non-Sequential = SUM]
usr_calls_t	The total number of user calls made by the session since it began. User calls are incremented whenever a parse, execute, or fetch is performed by the user. [Sequential = LST Non-Sequential = SUM]
usr_commits	The number of user commits performed by the session during the interval. User commits can be used to approximate a user transaction rate. [Sequential = SUM Non-Sequential = SUM]
usr_commits_t	The total number of user commits performed by the session since it began. User commits can be used to approximate a user transaction rate. [Sequential = LST Non-Sequential = SUM]
waits	The number of waits by the session during the interval [Sequential = SUM Non-Sequential = SUM]
waits_t	The total number of waits by the session since it began [Sequential = LST Non-Sequential = SUM]
waittime	The amount of time in seconds spent waiting by the session during the interval. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
waittime_t	The total amount of time in seconds spent waiting by the session since it began. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = LST Non-Sequential = SUM]

12.15. Session Wait Statistics

The Oracle.SessionWait table stores detailed information about current wait events by the session. The information is retrieved from the v\$session_event view within Oracle.

Table Field Hierarchy

Class:	Oracle
Subclass:	SessionWait
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Oracle/ <i>instancename</i>
TeamQuest Table Name:	Oracle.Session Wait
Open Table Name:	ORASESSIONWAIT
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avg_waittime	The average amount of time in seconds spent waiting for the event by the session during the interval. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = AVG Non-Sequential = AVG]
event	The name of the wait event [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
logon_time	The logon time of the session [Sequential = FST Non-Sequential = ID]
ora_user	The Oracle user name the session is using [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
serial_num	The session serial number [Sequential = ID Non-Sequential = ID]
sid	The session identifier [Sequential = ID Non-Sequential = ID]

System	<p>The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
timeouts	<p>The number of timeouts for the event by the session during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
timeouts_t	<p>The total number of timeouts for the event by the session since it began</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
waits	<p>The number of waits for the event by the session during the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
waits_t	<p>The total number of waits for the event by the session since it began</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
waittime	<p>The amount of time in seconds spent waiting for the event by the session during the interval. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
waittime_t	<p>The total amount of time in seconds spent waiting for the event since the session began. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>

12.16. System Parameters Statistics

The Oracle.SystemParameters table shows detailed information about the system parameters for an Oracle instance. The data is obtained from the v\$system_parameter.

Table Field Hierarchy

Class:	Oracle
Subclass:	SystemParameters
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.SystemParameters
Open Table Name:	ORASYSPARAMETERS
Collection interval:	N/A
Default retention:	6 months
Table type:	Event

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
description	The descriptive text about the parameter [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
isadjusted	Indicates that the rdbms adjusted the input value to a more suitable value. (For example, the parameter value should be prime, but the user input is a nonprime number, so the rdbms adjusted the value to the next prime number.) [Sequential = ID Non-Sequential = ID]
isdefault	Indicates the value assigned to the parameter is the default [Sequential = ID Non-Sequential = ID]
ismodified	Indicates how the parameter was modified. If an ALTER SESSION operation was performed, the value is MODIFIED. If an ALTER SYSTEM operation was performed (which causes all the values for the currently logged-in session to be modified), the value is SYS_MODIFIED. [Sequential = ID Non-Sequential = ID]
isses_modifiable	Indicates whether the parameter can be modified by ALTER SESSION [Sequential = ID Non-Sequential = ID]

issys_modifiable	Indicates whether the parameter can be modified by ALTER SYSTEM [Sequential = ID Non-Sequential = ID]
name	The parameter name [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
type	The parameter type. It can be one of the following: 1 = Boolean 2 = string 3 = integer [Sequential = ID Non-Sequential = ID]
value	The value assigned to the parameter [Sequential = ID Non-Sequential = ID]

12.17. System Statistics

The Oracle.SystemStats table shows detailed information about the system statistics that Oracle provides. The data is obtained from the v\$sysstat view. Most of the statistics in this table are also stored as aggregate values with the exception of the statistics that have a class type of 16, which is OS data. If a statistic in this table is a time field, the values are represented in seconds.

Table Field Hierarchy

Class:	Oracle
Subclass:	SystemStats
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.SystemStats
Open Table Name:	ORASYSSTATS
Collection interval:	N/A
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
class	A number representing a statistic class. The class numbers are defined as follows: 1 = User 2 = Redo 4 = Enqueue 8 = Cache 16 = OS 32 = Parallel Server 64 = SQL 128 = Debug [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
name	The name of the statistic [Sequential = ID Non-Sequential = ID]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
statnum	The statistic number. For a list of all the available statistics and a description of what each statistic collects, see the appropriate Oracle manual. [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
value	The difference between the statistic value during the interval [Sequential = SUM Non-Sequential = SUM]
value_t	The total number of the value [Sequential = LST Non-Sequential = SUM]

12.18. System Wait Event Statistics

The Oracle.SystemWait table stores detailed information about waits for system events within the Oracle data.

Table Field Hierarchy

Class:	Oracle
Subclass:	SystemWait
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.SystemWait
Open Table Name:	ORASYSWAIT
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avg_waittime	The average amount of time in seconds spent waiting for the event during the interval. The value is always zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = AVG Non-Sequential = AVG]

event	The name of the wait event [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
timeouts	The number of timeouts for the event during the interval [Sequential = SUM Non-Sequential = SUM]
timeouts_t	The total number of timeouts for the event since the instance started [Sequential = LST Non-Sequential = SUM]
waits	The number of waits for the event during the interval [Sequential = SUM Non-Sequential = SUM]
waits_t	The total number of waits for the event since the instance started [Sequential = LST Non-Sequential = SUM]
waittime	The amount of time in seconds spent waiting for the event during the interval. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = SUM Non-Sequential = SUM]
waittime_t	The amount of time in seconds spent waiting for the event since the instance started. This value is zero when the Oracle timed_statistics initialization parameter is set to FALSE. [Sequential = LST Non-Sequential = SUM]

12.19. Top SQL Cursors Statistics

The Oracle.TopSQL table stores snapshots of the top SQL Cursors for a given sample. This allows you to pinpoint possible query problems that result in poor performance on the system. You can select the number of rows returned and the criteria used for returning those rows. The information is obtained from the v\$sql view within Oracle. For more information on configuring the collection of top SQL cursor statistics, see the section on configuring an Oracle Agent in the *TeamQuest Performance Software Administration Guide*.

Table Field Hierarchy

Class:	Oracle
Subclass:	TopSQL
IT Resource Name:	/TeamQuest/System/systemname/Oracle/instancename
TeamQuest Table Name:	Oracle.TopSQL
Open Table Name:	ORATOPSQL
Collection interval:	24 hours (default)
Default retention:	30 days
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
buffer_gets	The number of buffer gets for the cursor during the interval [Sequential = SUM Non-Sequential = SUM]
buffer_gets_t	The total number of buffer gets for the cursor since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
cpu_time	The CPU time in seconds used by the cursor for parsing, executing, or fetching during the interval [Sequential = SUM Non-Sequential = SUM]
cpu_time_t	The total CPU time in seconds used by the cursor for parsing, executing, or fetching since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
disk_reads	The number of disk reads for the cursor during the interval [Sequential = SUM Non-Sequential = SUM]
disk_reads_t	The total number of disk reads for the cursor since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
elapsed_time	The elapsed time in seconds used by the cursor for parsing, executing, or fetching during the interval [Sequential = SUM Non-Sequential = SUM]

elapsed_time_t	The total elapsed time in seconds used by the cursor for parsing, executing, or fetching since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
executions	The number of executions that took place on the object during the interval [Sequential = SUM Non-Sequential = SUM]
executions_t	The total number of executions that took place on the object since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
first_load_time	The timestamp of the parent creation time [Sequential = FST Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained. This field is limited to 16 characters. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
invalidations	The number of times the cursor has been invalidated during the interval [Sequential = SUM Non-Sequential = SUM]
invalidations_t	The total number of times the cursor has been invalidated since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
loaded_versions	The number of child and parent cursors with a loaded context heap for the library cache entry [Sequential = LST Non-Sequential = SUM]
loads	The number of times the object was loaded or reloaded during the interval [Sequential = SUM Non-Sequential = SUM]
loads_t	The total number of times the object was loaded or reloaded since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
module	The name of the module that was executing at the time the SQL statement was first parsed as set by calling <code>dbms_applications_info.set_module</code> [Sequential = ID Non-Sequential = ID]
optimizer_mode	The mode under which the SQL statement is executed [Sequential = ID Non-Sequential = ID]
parse_calls	The number of parse calls for the cursor during the interval [Sequential = SUM Non-Sequential = SUM]
parse_calls_t	The total number of parse calls for the cursor since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
parsing_user_id	The user id of the user who originally built the cursor [Sequential = ID Non-Sequential = ID]

persistent_mem_KB	The fixed amount of memory in kilobytes (KB) used during the lifetime of the cursor [Sequential = LST Non-Sequential = SUM]
rows_processed	The number of rows the parsed SQL statement returned during the interval [Sequential = SUM Non-Sequential = SUM]
rows_processed_t	The total number of rows the parsed SQL statement returned since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
runtime_mem_KB	The fixed amount of memory in kilobytes (KB) required during the execution of the cursor [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
sharable_mem_KB	The amount of shared memory in kilobytes (KB) used by the cursor [Sequential = LST Non-Sequential = SUM]
sorts	The number of sort operations that was done for the cursor during the interval [Sequential = SUM Non-Sequential = SUM]
sorts_t	The total number of sort operations that was done for the cursor since it was present in the library cache [Sequential = LST Non-Sequential = SUM]
sql_address	The address of the handle to the parent for the cursor [Sequential = ID Non-Sequential = ID]
sql_hash_value	The hash value of the parent statement in the library cache [Sequential = ID Non-Sequential = ID]
sql_text	The SQL text for the current cursor [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
version_count	The number of cursors that are present in the library cache for this sql_address and sql_hash_value [Sequential = LST Non-Sequential = SUM]

Section 13

Oracle Solaris Systems

Statistics for Oracle Solaris systems are collected by the TeamQuest collection agents.

This section contains a listing of the statistics collected for the system:

- System Activity Statistics (see 13.1)
- Disk Space Statistics (see 13.2)
- Network Statistics (see 13.3)
- Workload Statistics (see 13.4)
- Process Statistics (see 13.5)
- Project Statistics (see 13.6)
- Hardware Inventory Statistics (see 13.7)
- System Log Statistics (see 13.8)
- General Log Statistics (see 13.9)
- TeamQuest Log Statistics (see 13.10)
- Derived Statistics (see 13.11)
- Zone Statistics (see 13.12)
- Processor Set Statistics (see 13.13)
- Resource Pool Statistics (see 13.14)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV =Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

13.1. System Activity Statistics

The System Activity Agent (**tqbsp**) is used to collect a wide variety of important system statistics. Major resources monitored by this agent include processors, memory, disks, and the operating system kernel.

Notes:

- *Two naming conventions are available for Block Device resources. By default, the resource names are in the form of c0t0d0s0 for local disks, rmt/0 for tapes, and server:/export for NFS Client mount points. If you would like to use the alternate names that are in the form of sd0 for disks, st0 for tapes, or nfs0 for NFS, change the Descriptive Disk Names setting in the System Activity Agent configuration file to OFF.*
- *By default, performance data on disk partitions is not stored. To collect performance data on disk partitions, change the Disk Partitions setting in the System Activity Agent configuration file to ON.*
- *By default, performance data NFS client mount point is not stored. To collect data on NFS client mount point, change the NFS Mounts setting in the System Activity Agent Configuration file to ON.*
- *For information on how to change the System Activity Agent configuration settings using TeamQuest Manager, see the TeamQuest Performance Software Administration Guide.*

Special Processing When Using Sequential Consolidation Method

Special processing occurs when certain records in the Block Device.by Device table are consolidated using the Sequential consolidation method. The following formulas are used to calculate the %busy, Actual_Interval, and record_count statistic values:

%busy

The %busy field uses a new consolidation method that uses the following formula to produce the consolidated %busy value:

$$\text{consolidated \%busy} = \%busy * \text{record_count} * \text{Actual_Interval}$$

At the end of the aggregation processing step after multiple records have been combined together to produce a single consolidated record, the %busy field contains the consolidated %busy value.

An additional processing step is performed using the following formula to produce a final %busy value that is stored into the consolidated record:

$$\%busy = \frac{\text{consolidated \%busy}}{\text{record_count} * \text{Interval}}$$

Note: The record_count field value used in the above formula must have already been generated using the record_count formula.

Block Device.by Device table records that have been stored by previous levels of TeamQuest collection agents do not contain the record_count field. For these records, a value of 1 is assumed for the record_count value.

Actual_Interval

For consolidated records (both reduced and not reduced), the Actual_Interval field should contain the Interval value at the end of the aggregation processing step.

record_count

The record_count field value is updated at the end of the aggregation processing step using the following formula:

$$\text{record_count} = \frac{(\sum(\text{record_count} * \text{Actual_Interval})) + (\text{Interval} - \sum \text{Actual_Interval})}{\text{Interval}}$$

The first part of the numerator is a summation of the record_count value multiplied by the Actual_Interval value across all of the consolidated records. This value should exist in the record_count field at the end of the aggregation processing step since the record_count field is weighted by Actual_Interval.

The Actual_Interval value used in the second part of the numerator is a summation of all the Actual_Interval values of the consolidated records. This value should exist in the Actual_Interval field at the end of the aggregation processing step since Actual_Interval is a summation.

The Interval value in the formula is the time range in seconds that the final consolidated record represents. For example, if 5-minute records are generated, the Interval value is calculated as 5 multiplied by 60.

Parameter Hierarchy

Class:	Block Device
Subclass:	by Device
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Block Device.by Device
Open Table Name:	BLKDEVBYDEVICE
Resource:	disk0, disk1, ...
Statistic Name:	
%busy	The percentage of time this device was servicing a transfer request [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-util.rpt
actq_avwait*	The average run queue wait time in milliseconds
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgresp*	The average response time of an I/O on a device. Calculated as avwait + avserv

avque	The average number of requests outstanding [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/io/dsk-q.rpt
avserv	The average time in milliseconds to service each transfer request (includes seek, rotation latency, and data transfer times) for the device [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-time.rpt
await	The average time in milliseconds that transfer requests are idle in the queue while the queue is occupied [Sequential = AVG Non-Sequential = AVG] View Report: /report/hp-ux/sys-act/io/dsk-time.rpt
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
IO_intensity*	The activity of an I/O device. This is the product of the I/O response time in milliseconds and the I/O transfer rate in I/Os per second. This is proportional to the average queue length (the number of I/O requests waiting or in progress at the I/O device).
Kbytes/s	The rate at which data is transferred in kilobytes per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/hp-ux/sys-act/io/dsk-xfer.rpt /report/hp-ux/sys-act/io/top-dsk.rpt
record_count	The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record. [Sequential = AVG Non-Sequential = SUM]
reduction_name	The name of reduction rule [Sequential = ID Non-Sequential = ID]
reduction_source	The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H. [Sequential = ID Non-Sequential = ID]
transfers/s	The number of physical transfers to and from the disk per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/hp-ux/sys-act/io/dsk-xfer.rpt
waitq_await*	The average wait queue wait time in milliseconds

Class:	Block Device
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	Block Device.Summary
Open Table Name:	BLKDEVSUM
Statistic Name:	
transfers/s	The number of physical transfers to and from all of the devices per second [Sequential = AVG Non-Sequential = SUM]
Class:	CPU
Subclass:	by Processor
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.by Processor
Open Table Name:	CPUBYPROC
Resource:	cpu0, cpu1, ...
Statistic Name:	
%idle	The percentage of CPU time spent idle for the CPU [Sequential = AVG Non-Sequential = AVG] View Reports: /report/solaris/sys-act/cpu/mpstat.rpt /report/solaris/sys-act/cpu/per-cpu.rpt
%sys	The percentage of CPU time spent in system mode for the CPU [Sequential = AVG Non-Sequential = AVG] View Reports: /report/solaris/sys-act/cpu/mpstat.rpt /report/solaris/sys-act/cpu/per-cpu.rpt
%usr	The percentage of CPU time spent in user mode for the CPU [Sequential = AVG Non-Sequential = AVG] View Reports: /report/solaris/sys-act/cpu/mpstat.rpt /report/solaris/sys-act/cpu/per-cpu.rpt
%wio	The percentage of CPU time spent idle while some process is waiting for I/O completion for the CPU. For Solaris level 10 and later, this statistic is no longer available and will appear as 0. %wio is reported as part of the %idle statistic. [Sequential = AVG Non-Sequential = AVG] View Reports: /report/solaris/sys-act/cpu/mpstat.rpt /report/solaris/sys-act/cpu/per-cpu.rpt
csw/s	The number of process switches per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt

icsw/s	The number of involuntary context switches per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
intr/s	The number of interrupts per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
ithr/s	The number of interrupts as threads (not counting clock interrupts) per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
migr/s	The number of thread migrations (to another CPU) per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
minf/s	The number of minor page faults per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
mjfl/s	The number of major page faults per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
smtx/s	The number of spins on mutexes (lock not acquired on first try) per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
srw/s	The number of spins on reader/writer locks (lock not acquired on first try) per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
syscl/s	The number of system calls per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt
xcal/s	The number of inter-processor cross calls per second for the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/cpu/mpstat.rpt

Table Field Hierarchy

Class:	CPU
Subclass:	RelativePerformance
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.RelativePerformance
Open Table Name:	CPURELPERF
Collection interval:	1 minute
Default retentions:	1 month
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Type	Description
Actual_Interval	Long	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
cpu_relative_performance	Real	The relative performance of the CPU on a common scale [Sequential = AVG Non-Sequential = SUM]
Interval	Long	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
rel_unused	Real	The amount of CPU resources not used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
rel_used	Real	The amount of CPU resources used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
System	String	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	Integer	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

Class:	CPU
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /CPU
TeamQuest Table Name:	CPU.Summary
Open Table Name:	CPUSUM
Statistic Name:	
%busy	The percentage of time the CPU was not idle [Sequential = AVG Non-Sequential = AVG]
%idle	The percentage of total CPU time spent idle while no processes are waiting for I/O completion [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/sys-act/cpu/cpu-util.rpt
%sys	The percentage of total CPU time spent in system mode [Sequential = AVG Non-Sequential = AVG] View Reports: /report/solaris/rules/cpu_mutex.rpt /report/solaris/sys-act/cpu/cpu-util.rpt
%usr	The percentage of total CPU time spent in user mode [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/sys-act/cpu/cpu-util.rpt
%wio	The percentage of total CPU time spent idle while some process is waiting for I/O completion. For Solaris level 10 and later, this statistic is no longer available and will appear as 0. %wio is reported as part of the %idle statistic. [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/sys-act/cpu/cpu-util.rpt
online_cpus	The number of CPUs that were online at the end of the sampling interval [Sequential = LST Non-Sequential = SUM]

Class:	Kernel
Subclass:	Buffers
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.Buffers
Open Table Name:	KNLBUFFS
Statistic Name:	
%rcache	The percentage of logical reads satisfied from the buffer cache [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/sys-act/kernel/bufc-hit.rpt
%wcache	The percentage of logical writes satisfied from the buffer cache [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/sys-act/kernel/bufc-hit.rpt
bread/s	The number of reads per second from devices into the buffer cache [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/bufc-xfr.rpt
bwrit/s	The number of writes per second from the buffer cache to devices [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/bufc-xfr.rpt
hitrate	The buffer cache hit ratio [Sequential = AVG Non-Sequential = AVG]
hits/s	The rate of buffer cache hits [Sequential = AVG Non-Sequential = SUM]
lookups/s	The rate of buffer cache lookups [Sequential = AVG Non-Sequential = SUM]
lread/s	The number of reads per second from the buffer cache to a process [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/bufc-xfr.rpt
lwrit/s	The number of writes per second from a process to the buffer cache [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/bufc-xfr.rpt
pread/s	The number of physical read requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/phys-xfr.rpt
pwrit/s	The number of physical write requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/phys-xfr.rpt

Class:	Kernel
Subclass:	File Access
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.File Access
Open Table Name:	KNLFILEACCESS
Statistic Name:	
dirblk/s	The number of directory block reads issued per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/f-access.rpt
dnlchitrate	The directory name lookup cache hit ratio [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/rules/dnlc.rpt
dnlchits/s	The number of directory name lookup cache hits per second [Sequential = AVG Non-Sequential = SUM]
dnlcmissses/s	The number of directory name lookup cache misses per second [Sequential = AVG Non-Sequential = SUM]
dnlcrefs/s	The number of directory name lookups per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/rules/dnlc.rpt
icrefs/s	The number of i-node lookups per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/solaris/rules/inode_hit.rpt /report/solaris/rules/inode_steal.rpt
iget/s	The number of files located by i-node entries per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/f-access.rpt
ihitrate	The i-node cache hit ratio [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/rules/inode_hit.rpt
ihits/s	The number of i-node cache hits per second [Sequential = AVG Non-Sequential = SUM]
imisses/s	The number of i-node cache misses per second [Sequential = AVG Non-Sequential = SUM]

iprate	The rate at which UNIX File System (UFS) i-nodes with reusable pages associated with them are taken off the free list [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/rules/inode_steal.rpt
namei/s	The number of file system path searches per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/f-access.rpt
Class:	Kernel
Subclass:	IPC (inter process communication)
IT Resource Name:	/TeamQuest/System/systemname/Kernel
TeamQuest Table Name:	Kernel.IPC
Open Table Name:	KNLIPC
Statistic Name:	
msg/s	The number of message operations (send and receives) per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/msg-sema.rpt
sema/s	The number of semaphore operations per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/msg-sema.rpt
Class:	Kernel
Subclass:	Load Average
IT Resource Name:	/TeamQuest/System/systemname/Kernel
TeamQuest Table Name:	Kernel.Load Average
Open Table Name:	KNLLOADAVG
Statistic Name:	
1 min	The number of processes in the run queue averaged over the last 1 minute. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/load-avg.rpt

5 min	<p>The number of processes in the run queue averaged over the last 5 minutes. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/load-avg.rpt</p>
15 min	<p>The number of processes in the run queue averaged over the last 15 minutes. This count is taken at the end of the sampling interval. [Sequential = LST Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/load-avg.rpt</p>
Class:	Kernel
Subclass:	Memory
IT Resource Name:	/TeamQuest/System/systemname/Kernel
TeamQuest Table Name:	Kernel.Memory
Open Table Name:	KNLMEM
Statistic Name:	
failures	<p>This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = LST Non-Sequential = SUM]</p>
in_use	<p>This statistic is not available for the Oracle Database Server. The value is reported as <N/A>. [Sequential = LST Non-Sequential = SUM]</p>
Class:	Kernel
Subclass:	Paging
IT Resource Name:	/TeamQuest/System/systemname/Memory
TeamQuest Table Name:	Kernel.Paging
Open Table Name:	KNLPAGING
Statistic Name:	
%ufsipf	<p>The percentage of UNIX File System (UFS) i-nodes taken off the free list by iget, which had reusable pages associated with it [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/sys-act/kernel/paging2.rpt</p>
apfree/s	<p>The number of anonymous pages freed per second [Sequential = AVG Non-Sequential = SUM]</p>
appgin/s	<p>The number of anonymous pages paged-in per second [Sequential = AVG Non-Sequential = SUM]</p>
appgout/s	<p>The number of anonymous pages paged-out per second [Sequential = AVG Non-Sequential = SUM]</p>

atch/s	The number of page faults per second that are satisfied by reclaiming a page currently in memory (attaches per second) [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging1.rpt
epfree/s	The number of executable pages freed per second [Sequential = AVG Non-Sequential = SUM]
eppgin/s	The number of executable pages paged-in per second [Sequential = AVG Non-Sequential = SUM]
eppgout/s	The number of executable pages paged-out per second [Sequential = AVG Non-Sequential = SUM]
fpfree/s	The number of file system pages freed per second [Sequential = AVG Non-Sequential = SUM]
fppgin/s	The number of file system pages paged-in per second [Sequential = AVG Non-Sequential = SUM]
fppgout/s	The number of file system pages paged-out per second [Sequential = AVG Non-Sequential = SUM]
handspreadpages	The maximum distance between clock hands in bytes. This value will be one-fourth of all memory up to a maximum of 8,192 bytes. [Sequential = AVG Non-Sequential = AVG]
pflts/s	The number of page faults per second from protection errors or copy-on-writes [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging1.rpt
pgfree/s	The number of pages per second placed on the free list by the page-stealing daemon [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging2.rpt
pgin/s	The number of page-in requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging1.rpt
pgout/s	The number of page-out requests per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging2.rpt
pgscan/s	The number of pages per second scanned by the page-stealing daemon [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging2.rpt
ppgin/s	The number of pages paged-in per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging1.rpt

ppgout/s	The number of pages paged-out per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging2.rpt
slock/s	The number of faults per second caused by software locks requiring physical I/O [Sequential = AVG Non-Sequential = SUM]
vflts/s	The number of address translation page faults per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/paging1.rpt
Class:	Kernel
Subclass:	Queues
IT Resource Name:	/TeamQuest/System/systemname/Kernel
TeamQuest Table Name:	Kernel.Queues
Open Table Name:	KNLQS
Statistic Name:	
%runocc	The percentage of time the run queue is occupied [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/sys-act/kernel/q-util.rpt
%swpocc	The percentage of time the swap queue is occupied [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/sys-act/kernel/q-util.rpt
avg_cpuq_sz	The average length of the run queue per CPU (a queue of processes in memory and runnable) [Sequential = AVG Non-Sequential = AVG]
avg_runq_sz	The average length of the run queue (a queue of processes in memory and runnable) [Sequential = AVG Non-Sequential = AVG]
avg_swpq_sz	The average length of the swap queue (a queue of processes swapped out and ready to run) [Sequential = AVG Non-Sequential = AVG]
blocked	The average number of processes blocked waiting for resources [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/blocked.rpt
cpuq_sz	The average length of the run queue per CPU (a queue of processes in memory and runnable) while the run queue is occupied [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/rules/cpu_power.rpt

pswch/s	<p>The number of process switches per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/p-switch.rpt</p>
runq_sz	<p>The average length of the run queue (a queue of processes in memory and runnable) while the run queue is occupied [Sequential = AVG Non-Sequential = AVG] View Reports: /report/solaris/sys-act/kernel/q-sizes.rpt /report/solaris/sys-act/kernel/runq.rpt</p>
smtx/s	<p>The CPU mutex rate per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/rules/cpu_mutex.rpt</p>
swpq_sz	<p>The average length of the swap queue (a queue of processes swapped out and ready to run) while the swap queue is occupied [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/q-sizes.rpt</p>
Class:	Kernel
Subclass:	Swapping
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Kernel.Swapping
Open Table Name:	KNLSWAPPING
Statistic Name:	
allocated	<p>The total amount of swap space in megabytes allocated for use as backing store. The count is taken at the end of the sampling interval. [Sequential = AVG Non-Sequential = SUM]</p>
available	<p>The total swap space in megabytes that is available for future reservation and allocation. The count is taken at the end of the sampling interval. [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/freeswap.rpt</p>
bswin/s	<p>The number of blocks transferred per second for swap-ins [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/swapamt.rpt</p>
bswot/s	<p>The number of blocks transferred per second for swap-outs [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/swapamt.rpt</p>
reserved	<p>The total amount of swap space in megabytes not allocated but claimed by memory mappings to be available for future use. The count is taken at the end of the sampling interval. [Sequential = AVG Non-Sequential = SUM]</p>

swpin/s	<p>The rate of swap-ins (transfers into memory per second)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/sys-act/kernel/swaprate.rpt</p>
swpot/s	<p>The rate of swap-outs (transfers from memory to the swap area per second)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/sys-act/kernel/swaprate.rpt</p>
used	<p>The total amount of swap space in megabytes that is either allocated or reserved. The count is taken at the end of the sampling interval.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Class:	Kernel
Subclass:	Tables
IT Resource Name:	/TeamQuest/System/systemname/Kernel
TeamQuest Table Name:	Kernel.Tables
Open Table Name:	KNLTABS
Statistic Name:	
inod_sz	<p>The number of entries currently used in the i-node table. This count is taken at the end of the sampling interval.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/solaris/sys-act/kernel/tbl-size.rpt</p>
proc_maxsz	<p>The maximum possible entries in the process table. This count is taken at the end of the sampling interval.</p> <p>[Sequential = MAX Non-Sequential = MAX]</p> <p>View Report:</p> <p>/report/solaris/sys-act/kernel/tbl-size.rpt</p>
proc_sz	<p>The number of entries currently being used in the process table. This count is taken at the end of the sampling interval.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/solaris/sys-act/kernel/tbl-size.rpt</p>

Class:	Kernel
Subclass:	TTY
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	Kernel.TTY
Open Table Name:	KNLTTY
Statistic Name:	
canch/s	The number of input characters per second processed by canon (canonical queue) [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/tty-xfer.rpt
mdmin/s	The number of modem interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/tty-intr.rpt
outch/s	The number of output characters transferred per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/tty-xfer.rpt
rawch/s	The number of input characters per second transferred in raw mode [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/tty-xfer.rpt
rcvin/s	The number of receiver hardware interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/tty-intr.rpt
xmtin/s	The number of transmitter hardware interrupts per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/kernel/tty-intr.rpt

Class:	Memory
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Memory
Open Table Name:	MEM
Statistic Name:	
zfs-arc-sz	<p>The total amount of memory in megabytes used by Adaptive Replacement Cache at the end of the interval. This statistic is available only when the Zettabyte File System (ZFS) is used on the system. This memory is available to user processes and is returned to freemem as required.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
freemem	<p>The average amount of memory in megabytes available to user processes not including memory allocated to ZFS.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/solaris/sys-act/memory/freemem.rpt</p> <p>/report/solaris/sys-act/memory/memory.rpt</p>
physmem	<p>The total amount of physical memory in megabytes at the end of the interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/sys-act/memory/memory.rpt</p>
Class:	Swap Space
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Memory
TeamQuest Table Name:	Swap Space.Summary
Open Table Name:	SWAPSPACESUM
Statistic Name:	
free	<p>The amount of free swap space in megabytes at the end of the interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/solaris/rules/swap_space.rpt</p> <p>/report/solaris/sys-act/swap/swpspc.rpt</p>
total	<p>The total amount of swap space in megabytes at the end of the interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/sys-act/swap/swpspc.rpt</p>

Class:	System Call
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Kernel
TeamQuest Table Name:	System Call.Summary
Open Table Name:	SYSCALLSUM
Statistic Name:	
exec/s	The number of exec system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/syscall/imp-scal.rpt
fork/s	The number of fork and vfork system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/syscall/imp-scal.rpt
rchar/s	The number of characters transferred by read system calls in the interval in bytes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/syscall/scal-xfr.rpt
scall/s	The total number of system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/syscall/imp-scal.rpt
sread/s	The number of read system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/syscall/imp-scal.rpt
swrit/s	The number of write system calls per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/syscall/imp-scal.rpt
wchar/s	The number of characters transferred by write system calls in the interval in bytes per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/sys-act/syscall/scal-xfr.rpt

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
bsp interval	The number of seconds elapsed between two data samples of the System Activity Agent [Sequential = SUM Non-Sequential = ID]
tqbsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqbsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

Parameter Hierarchy

Class:	ZFS
Subclass:	ARC Detail
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	ZFS.ARC Detail
Open Table Name:	ZFSARCDETAIL
Resource:	N/A
Statistic Name:	
demand_hit_pct	The percentage of Adaptive Replacement Cache (ARC) demand reads that resulted in cache hits. Demand reads are all reads not originated by a Zettabyte File System (ZFS) predictive algorithm, such as ZFS prefetch. [Sequential = AVG Non-Sequential = SUM]
demand_hits_per_sec	The number of ARC demand hits per second [Sequential = AVG Non-Sequential = SUM]
demand_miss_pct	The percentage of ARC demand reads that resulted in cache misses [Sequential = AVG Non-Sequential = SUM]
demand_misses_per_sec	The number of ARC demand misses per second [Sequential = AVG Non-Sequential = SUM]
demand_reads_per_sec	The number of ARC demand reads per second [Sequential = AVG Non-Sequential = SUM]
evict_skips_per_sec	The number of ARC evictions that were skipped (not performed) per second. Prefetched data in the ARC has a minimum life span, forcing the ARC eviction to skip the prefetched data. [Sequential = AVG Non-Sequential = SUM]
metadata_current	The current size of the metadata cache in megabytes stored in the ARC. The metadata cache is used to hold file system metadata. Metadata includes file, directory, and volume information. [Sequential = AVG Non-Sequential = SUM]
metadata_hit_pct	The percentage of ARC metadata reads that resulted in cache hits [Sequential = AVG Non-Sequential = SUM]
metadata_hits_per_sec	The number of ARC metadata hits per second. Metadata requests relate to file system or volume metadata held within the cache. Metadata includes file, directory, and volume information. [Sequential = AVG Non-Sequential = SUM]
metadata_limit	The tunable metadata limit in megabytes for the metadata cache stored in the ARC. A value of zero indicates that the tunable value is not set and that the actual limit is determined by the operating system. [Sequential = AVG Non-Sequential = SUM]
metadata_maximum	The maximum size of the metadata cache in megabytes stored in the ARC [Sequential = AVG Non-Sequential = SUM]

metadata_miss_pct	The percentage of ARC metadata reads that resulted in cache misses [Sequential = AVG Non-Sequential = SUM]
metadata_misses_per_sec	The number of ARC metadata misses per second [Sequential = AVG Non-Sequential = SUM]
metadata_reads_per_sec	The number of ARC metadata reads per second. Calculated as hits + misses. [Sequential = AVG Non-Sequential = SUM]
mfu_ghost_hits_per_sec	The number of Most Frequently Used (MFU) ARC ghost hits per second. Ghost hits occur when requested data is found in the MFU ghost cache. The MFU ghost cache contains metadata on items recently evicted from the MFU cache and is used to dynamically tune ARC caching. [Sequential = AVG Non-Sequential = SUM]
mfu_hits_per_sec	The number of Most Frequently Used (MFU) ARC hits per second. The MFU cache is internal to the ARC and is used with the Most Recently Used (MRU) cache to balance overall cache performance. [Sequential = AVG Non-Sequential = SUM]
mru_ghost_hits_per_sec	The number of Most Recently Used (MRU) ARC ghost hits per second. Ghost hits occur when requested data is found in the MRU ghost cache. The MRU ghost cache contains metadata on items recently evicted from the MRU cache and is used to dynamically tune ARC caching. [Sequential = AVG Non-Sequential = SUM]
mru_hits_per_sec	The number of Most Recently Used (MRU) ARC hits per second. The MRU cache is internal to the ARC and is used with the Most Frequently Used (MFU) cache to balance overall ARC performance. [Sequential = AVG Non-Sequential = SUM]
mutex_misses_per_sec	The number of ARC mutex misses per second. A mutex miss occurs when ARC data cannot be evicted due to the inability of the eviction process to acquire a lock on the data. This inability is caused by processing on the cached data by another part of the overall ARC process. [Sequential = AVG Non-Sequential = SUM]
prefetch_hit_pct	The percentage of ARC prefetch reads that resulted in cache hits. Prefetch requests occur when the operating system has determined that sequential I/O operations are occurring and the prefetching of data will result in improved cache performance. These requests originate from the operating system and do not directly reflect user-generated I/O operations. [Sequential = AVG Non-Sequential = SUM]
prefetch_hits_per_sec	The number of ARC prefetch hits per second [Sequential = AVG Non-Sequential = SUM]
prefetch_miss_pct	The percentage of ARC prefetch reads that resulted in cache misses [Sequential = AVG Non-Sequential = SUM]

prefetch_misses_per_sec	The number of ARC prefetch misses per second [Sequential = AVG Non-Sequential = SUM]
prefetch_reads_per_sec	The number of ARC prefetch reads per second. Calculated as hits + misses. [Sequential = AVG Non-Sequential = SUM]
recycle_misses_per_sec	The number of ARC recycle misses per second. Cache recycles occur when an allocated block of data is reused, rather than freeing an evicted block of data. This recycling avoids the overhead of freeing and immediately reallocating a block of memory of an identical size. Recycle misses occur when cache evictions are requested but no cache items of the appropriate size are found for recycling. [Sequential = AVG Non-Sequential = SUM]

Parameter Hierarchy

Class:	ZFS
Subclass:	ARC Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	ZFS.ARC Summary
Open Table Name:	ZFSARCSUMMARY
Resource:	N/A
Statistic Name:	
hit_pct	The percentage of Adaptive Replacement Cache (ARC) reads that resulted in cache hits [Sequential = AVG Non-Sequential = SUM]
hits_per_sec	The number of ARC hits per second. This value is the sum of demand hits and prefetch hits. [Sequential = AVG Non-Sequential = SUM]
miss_pct	The percentage of ARC reads that resulted in cache misses [Sequential = AVG Non-Sequential = SUM]
misses_per_sec	The number of ARC misses per second. This value is the sum of the demand misses and prefetch misses. [Sequential = AVG Non-Sequential = SUM]
reads_per_sec	The number of ARC reads per second. This value is the sum of total cache hits and misses per second. [Sequential = AVG Non-Sequential = SUM]
size	The current size of the ARC in megabytes (MB) [Sequential = AVG Non-Sequential = SUM]
target_size	The target size of the ARC. This value is the ideal size of the cache as determined by the operating system. [Sequential = AVG Non-Sequential = SUM]

Parameter Hierarchy

Class:	ZFS
Subclass:	L2 ARC
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	ZFS.L2 ARC
Open Table Name:	ZFSL2ARC
Resource:	N/A
Statistic Name:	
hit_pct	The percentage of Level 2 (L2) Adaptive Replacement Cache (ARC) reads that resulted in cache hits [Sequential = AVG Non-Sequential = SUM]
hits_per_sec	The number of L2 ARC hits per second [Sequential = AVG Non-Sequential = SUM]
miss_pct	The percentage of L2 ARC reads that resulted in cache misses [Sequential = AVG Non-Sequential = SUM]
misses_per_sec	The number of L2 ARC misses per second [Sequential = AVG Non-Sequential = SUM]
rbytes_per_sec	The number of bytes read from the L2 ARC per second [Sequential = AVG Non-Sequential = SUM]
reads_per_sec	The number of L2 ARC reads per second. This value is equal to the sum of total L2 cache hits and misses per second. [Sequential = AVG Non-Sequential = SUM]
size	The current size of the L2 ARC in megabytes (MB) [Sequential = AVG Non-Sequential = SUM]

Parameter Hierarchy

Class:	ZFS
Subclass:	Virtual Device Cache
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	ZFS.Virtual Device Cache
Open Table Name:	ZFSVDEVCACHE
Resource:	N/A
Statistic Name:	
delegation_pct	The percentage of virtual device cache reads that resulted in cache hits on delegated cache data [Sequential = AVG Non-Sequential = SUM]
delegations_per_sec	The number of virtual device cache reads per second that resulted in cache hits on delegated cache data. A delegation is a token that a disk device grants the virtual device cache to ensure that the cache can serve read requests without the need for validating the data with a disk device. [Sequential = AVG Non-Sequential = SUM]

hit_pct	The percentage of virtual device cache reads that resulted in cache hits [Sequential = AVG Non-Sequential = SUM]
hits_per_sec	The number of virtual device cache hits per second [Sequential = AVG Non-Sequential = SUM]
miss_pct	The percentage of virtual device cache reads that resulted in cache misses [Sequential = AVG Non-Sequential = SUM]
misses_per_sec	The number of virtual device cache misses per second [Sequential = AVG Non-Sequential = SUM]
reads_per_sec	The number of virtual device cache reads per second. Calculated as hits_per_sec + misses_per_sec + delegations_per_sec. [Sequential = AVG Non-Sequential = SUM]

Parameter Hierarchy

Class:	ZFS
Subclass:	ZFetch
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Disk
TeamQuest Table Name:	ZFS.ZFetch
Open Table Name:	ZFSZFETCH
Resource:	N/A
Statistic Name:	
bogus_streams_per_sec	The number of invalid streams found per second [Sequential = AVG Non-Sequential = SUM]
colinear_hits_per_sec	The number of times per second that two linear sequences of accesses were combined into one strided sequence [Sequential = AVG Non-Sequential = SUM]
colinear_misses_per_sec	The number of times per second that two linear sequences of accesses were not combined into one strided sequence [Sequential = AVG Non-Sequential = SUM]
hits_per_sec	The number of cache hits per second [Sequential = AVG Non-Sequential = SUM]
misses_per_sec	The number of cache misses per second [Sequential = AVG Non-Sequential = SUM]
reclaim_failures_per_sec	The number of times per second that an existing stream could not be reclaimed for a new use [Sequential = AVG Non-Sequential = SUM]
reclaim_successes_per_sec	The number of times per second that an existing stream was successfully reclaimed for a new use [Sequential = AVG Non-Sequential = SUM]
streams_noresets_per_sec	The number of times per second that an existing stream was reused without being reset [Sequential = AVG Non-Sequential = SUM]

streams_resets_per_sec	The number of times per second that an existing stream was reset for a new use [Sequential = AVG Non-Sequential = SUM]
stride_hits_per_sec	The number of cache hits per second due to a strided prefetch [Sequential = AVG Non-Sequential = SUM]
stride_misses_per_sec	The number of cache misses per second during a strided prefetch [Sequential = AVG Non-Sequential = SUM]

13.2. Disk Space Statistics

The Disk Space Agent (**tqdsp**) tracks the disk space usage of locally mounted file systems. However, any data about Network File System (NFS) mounted file systems are not collected by the agent because the loss of the network connection to the NFS may result in hung processes.

Note: File system names longer than 51 characters will be truncated.

Parameter Hierarchy

Class:	Disk Space
Subclass:	by File System
IT Resource Name:	/TeamQuest/System/systemname/Disk
TeamQuest Table Name:	Disk Space.by File System
Open Table Name:	DISKSPACEBYFILESYS
Resource:	file system1, file system2, ...
Statistic Name:	
%inodes free*	The percentage of i-nodes available (not in use) on the file system at the end of the interval View Report: /report/solaris/dskspace/total/pct-inod.rpt
%inodes used*	The percentage of i-nodes in use on the file system at the end of the interval View Report: /report/solaris/dskspace/total/pct-inod.rpt
%space free*	The percentage of total space available (not in use) on the file system at the end of the interval View Reports: /report/solaris/dskspace/total/pctspace.rpt /report/solaris/dskspace/total/low-ones.rpt
%space used*	The percentage of total space in use on the file system at the end of the interval View Reports: /report/solaris/dskspace/total/pctspace.rpt /report/solaris/dskspace/total/fullest.rpt

%user space free*	<p>The percentage of total user space available (not in use) on the file system at the end of the interval</p> <p>View Reports:</p> <p>/report/solaris/dskspace/user/pctspace.rpt</p> <p>/report/solaris/dskspace/user/low-ones.rpt</p>
%user space used*	<p>The percentage of total user space in use on the file system at the end of the interval</p> <p>View Reports:</p> <p>/report/solaris/dskspace/user/pctspace.rpt</p> <p>/report/solaris/dskspace/user/fullest.rpt</p>
capacity	<p>The percentage of total space in use on the file system at the end of the interval</p> <p>[Sequential = LST Non-Sequential = AVG]</p>
free (Mb)	<p>The amount of space available (not in use) on the file system in megabytes (Mb) at the end of the interval including the space held back from normal users</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/solaris/dskspace/total/dskspace.rpt</p> <p>/report/solaris/dskspace/total/low-ones.rpt</p>
free inodes	<p>The number of available (not in use) i-nodes on the file system at the end of the interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/dskspace/total/i-nodes.rpt</p>
total (Mb)	<p>The total (used + available) amount of space on the file system in megabytes (Mb) at the end of the interval including the space held back from normal users</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/dskspace/total/dskspace.rpt</p>
total inodes	<p>The total (used + available) number of i-nodes on the file system at the end of the interval</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/dskspace/total/i-nodes.rpt</p>
user free (Mb)	<p>The amount of space available (not in use) on the file system in megabytes (Mb) at the end of the interval not including the space held back from normal users</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Reports:</p> <p>/report/solaris/dskspace/user/dskspace.rpt</p> <p>/report/solaris/dskspace/user/low-ones.rpt</p>
user total (Mb)*	<p>The total (used + available) amount of space on the file system in megabytes (Mb) at the end of the interval not including the space held back from normal users</p> <p>View Report:</p> <p>/report/solaris/dskspace/user/dskspace.rpt</p>

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
dsp interval	The number of seconds elapsed between two data samples of the Disk Space Agent [Sequential = SUM Non-Sequential = ID]
tqdsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqdsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

13.3. Network Statistics

The Network Agent (**tqbnp**) collects data about the interaction of the system with the network. The types of data collected by the agent include network interfaces, NFS and Remote Protocol Call (RPC), and Transmission Control Protocol (TCP).

In NFS environments, the term *server* refers to the system that owns a file system and allows other systems on the network to access the files on the file system. A *client* is a system that mounts a nonlocal file system such as NFS and accesses files from the NFS mounted file system. Any given system could be a server to some file systems and a client to other file systems. It is not uncommon to have systems that export several file systems for other systems to use and mount several other network file systems. All the NFS statistics collected by this agent show the activity of NFS clients and servers that exist on the system where the agent is running.

Parameter Hierarchy

Class:	Network
Subclass:	by Interface
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	Network.by Interface
Open Table Name:	NETBYINTERFACE
Resource:	interface0, interface1, ...
Statistic Name:	
brdrcv/s	The number of network interface broadcast bytes received per second [Sequential = AVG Non-Sequential = SUM]
brdctxmt/s	The number of network interface broadcast bytes transmitted per second [Sequential = AVG Non-Sequential = SUM]
collisions/s	The number of network collisions per second on Carrier Sense Multiple Access (CSMA) interfaces. This statistic is not available on all types of network interfaces. The values for this statistic will be zero for interface types that do not support this measurement. [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/net-errs.rpt
defer/s	The number of network interface packets that could not be immediately sent per second [Sequential = AVG Non-Sequential = SUM]
ifspeed	The number of network interface line speed in megabits per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/net-bits.rpt

in errors/s	The number of network input errors per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/solaris/network/net-errs.rpt /report/solaris/rules/network_errors.rpt
in packets/s	The number of network input packets per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/net-pkts.rpt
Mbits/s	The network interface bit count in megabits per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/net-bits.rpt
multircv/s	The number of network interface multicast bytes received per second [Sequential = AVG Non-Sequential = SUM]
multixmt/s	The number of network interface multicast bytes transmitted per second [Sequential = AVG Non-Sequential = SUM]
nocanput/s	The number of network interface packets dropped by IP per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/rules/network_transmit.rpt
norcvbuf/s	The number of network interface buffer receive allocation failures per second [Sequential = AVG Non-Sequential = SUM]
noxmtbuf/s	The number of network interface transmit buffer allocation failures per second [Sequential = AVG Non-Sequential = SUM]
obytes/s	The number of network interface bytes output per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/net-bytes.rpt
out errors/s	The number of network output errors per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/solaris/network/net-errs.rpt /report/solaris/rules/network_errors.rpt

out packets/s	<p>The number of network output packets per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/solaris/network/net-pkts.rpt /report/solaris/rules/network_collision.rpt /report/solaris/rules/network_defer.rpt</p>
rbytes/s	<p>The number of network interface bytes read per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/net-bytes.rpt</p>
Class:	Network
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/systemname/Network
TeamQuest Table Name:	Network.Summary
Open Table Name:	NETSUM
Statistic Name:	
errors/s	<p>The total number of network errors per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]</p>
in Kbytes/s	<p>The total number of network interface bytes input per second in kilobytes [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/sum-bytes.rpt</p>
in packets/s	<p>The total number of network input packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/net-sum.rpt</p>
out Kbytes/s	<p>The total number of network interface bytes output per second in kilobytes [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/sum-bytes.rpt</p>
out packets/s	<p>The total number of network output packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/net-sum.rpt</p>
total Kbytes/s	<p>The total number of network interface bytes input and output per second in kilobytes [Sequential = AVG Non-Sequential = SUM]</p>
total packets/s	<p>The total number of network packets per second for all network interfaces on the system [Sequential = AVG Non-Sequential = SUM]</p>

Class:	NFS
Subclass:	Client
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFS.Client
Open Table Name:	NFSCLI
Statistic Name:	
badcalls/s	<p>The total number of Network File System (NFS) calls per second rejected from the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/client.rpt</p>
calls/s	<p>The total number of NFS calls sent by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/client.rpt</p>
gets/s	<p>The total number of times per second an NFS client handle was received</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/client.rpt</p>
Class:	NFS
Subclass:	Server
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	NFS.Server
Open Table Name:	NFSSERV
Statistic Name:	
badcalls/s	<p>The total number of NFS calls per second rejected by the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/server.rpt</p>
calls/s	<p>The total number of NFS calls per second received by the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/server.rpt</p>

Class: NFSv2
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv2.Client
Open Table Name: NFSV2CLI
Statistic Name:
 calls/s* The number of NFS version 2 calls per second sent by the client
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/solaris/network/clnt-v2s.rpt

Class: NFSv2
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv2.Client
Open Table Name: NFSV2CLI
Resource: create, getattr, link, lookup, mkdir, null, read, readdir, readlink,
 remove, rename, rmdir, root, setattr, statfs, symlink, write, writecache
Statistic Name:
 reqs/s The number of NFS version 2 requests per second sent by the client
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/solaris/network/clnt-v2.rpt

Class: NFSv2
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv2.Server
Open Table Name: NFSV2SERV
Statistic Name:
 calls/s* The number of NFS version 2 calls per second received by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/solaris/network/srvr-v2s.rpt

Class: NFSv2
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv2.Server
Open Table Name: NFSV2SERV
Resource: create, getattr, link, lookup, mkdir, null, read, readdir, readlink, remove, rename, rmdir, root, setattr, statfs, symlink, write, writecache

Statistic Name:
 reqs/s
The number of NFS version 2 requests per second received by the server
[Sequential = AVG Non-Sequential = SUM]
View Report:
/report/solaris/network/srvr-v2.rpt

Class: NFSv3
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Client
Open Table Name: NFSV3CLI

Statistic Name:
 calls/s*
The number of NFS version 3 calls per second sent by the client
[Sequential = AVG Non-Sequential = SUM]
View Report:
/report/solaris/network/clnt-v3s.rpt

Class: NFSv3
Subclass: Client
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSV3.Client
Open Table Name: NFSV3CLI
Resource: access, commit, create, fsinfo, fsstat, getattr, link, lookup, mkdir, mknod, null, pathconf, read, readdir, readdir+, readlink, remove, rename, rmdir, setattr, symlink, write

Statistic Name:
 reqs/s
The number of NFS version 3 requests per second sent by the client
[Sequential = AVG Non-Sequential = SUM]
View Report:
/report/solaris/network/clnt-v3.rpt

Class: NFSv3
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Server
Open Table Name: NFSV3SERV
Statistic Name:
 calls/s* The number of NFS version 3 calls per second received by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/solaris/network/srvr-v3s.rpt

Class: NFSv3
Subclass: Server
IT Resource Name: /TeamQuest/System/*systemname*/Network
TeamQuest Table Name: NFSv3.Server
Open Table Name: NFSV3SERV
Resource: access, commit, create, fsinfo, fsstat, getattr, link, lookup, mkdir, mknod, null, pathconf, read, readdir, readdir+, readlink, remove, rename, rmdir, setattr, symlink, write
Statistic Name:
 reqs/s The number of NFS version 3 requests per second received by the server
 [Sequential = AVG Non-Sequential = SUM]
 View Report:
 /report/solaris/network/srvr-v3.rpt

Class:	RPC
Subclass:	Client.Connectionless
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	RPC.Client.Connectionless
Open Table Name:	RPCCLICONNLESS
Statistic Name:	
badcalls/s	<p>The number of connectionless RPC calls per second rejected by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/network/rpc/clnt-cl.rpt</p>
badxid/s	<p>The number of times per second a reply from a server was received that did not correspond to any outstanding connectionless RPC call</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/network/rpc/clnt-cl.rpt</p>
calls/s	<p>The total number of connectionless RPC calls per second sent by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/network/rpc/clnt-cl.rpt</p>
newcred/s	<p>The number of times per second connectionless RPC authentication information had to be refreshed by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/network/rpc/clnt-cl.rpt</p>
retrans/s	<p>The number of times per second a connectionless RPC call had to be retransmitted by the client due to a timeout while waiting for a reply from the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/network/rpc/clnt-cl.rpt</p>
timeout/s	<p>The number of times per second a connectionless RPC call timed out while waiting for a reply from the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/network/rpc/clnt-cl.rpt</p>

Class:	RPC
Subclass:	Client.Connection Oriented
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	RPC.Client.Connection Oriented
Open Table Name:	RPCCLICONNORIENTED
Statistic Name:	
badcalls/s	<p>The number of connection-oriented RPC calls per second rejected from the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/rpc/clnt-co.rpt</p>
badxid/s	<p>The number of times per second a reply from a server was received that did not correspond to any outstanding connection-oriented RPC call</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/rpc/clnt-co.rpt</p>
calls/s	<p>The total number of connection-oriented RPC calls per second sent by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/rpc/clnt-co.rpt</p>
newcred/s	<p>The number of times per second connection-oriented RPC authentication information had to be refreshed by the client</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/rpc/clnt-co.rpt</p>
timeout/s	<p>The number of times per second a connection-oriented RPC call timed out while waiting for a reply from the server</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/solaris/network/rpc/clnt-co.rpt</p>

Class:	RPC
Subclass:	Server.Connectionless
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	RPC.Server.Connectionless
Open Table Name:	RPCSERVCONNLESS
Statistic Name:	
badcalls/s	<p>The number of connectionless RPC calls per second rejected by the server. The sum of badlens/s and xdrcalls/s. [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-cl.rpt</p>
badlens/s	<p>The number of connectionless RPC calls per second received by the server with a length shorter than a minimum-sized RPC call [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-cl.rpt</p>
calls/s	<p>The number of connectionless RPC calls per second received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-cl.rpt</p>
nullrecvs/s	<p>The number of times per second a connectionless RPC call was not available when it was thought to be received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-cl.rpt</p>
xdrcalls/s	<p>The number of connectionless RPC calls per second by the server whose header could not be External Data Representation (XDR) decoded [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-cl.rpt</p>

Class:	RPC
Subclass:	Server.Connection Oriented
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	RPC.Server.Connection Oriented
Open Table Name:	RPCSERVCONNORIENTED
Statistic Name:	
badcalls/s	<p>The number of connection-oriented RPC calls per second rejected by the server. The sum of badlens/s and xdrcalls/s. [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-co.rpt</p>
badlens/s	<p>The number of connection-oriented RPC calls per second received by the server with a length shorter than a minimum-sized RPC call [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-co.rpt</p>
calls/s	<p>The number of connection-oriented RPC calls per second received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-co.rpt</p>
nullrecvs/s	<p>The number of times per second a connection-oriented RPC call was not available when it was thought to be received by the server [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-co.rpt</p>
xdrcalls/s	<p>The number of connection-oriented RPC calls per second by the server whose header could not be XDR decoded [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/network/rpc/srvr-co.rpt</p>

Class:	TCP
Subclass:	N/A
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Network
TeamQuest Table Name:	TCP
Open Table Name:	TCP
Statistic Name:	
ftp conn	<p>The number of FTP connections at the end of an interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/network/tcp/connect.rpt</p>
http conn	<p>The number of HTTP connections at the end of the interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/network/tcp/connect.rpt</p>
other conn	<p>The number of other connections at the end of an interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/network/tcp/connect.rpt</p>
rlogin conn	<p>The number of rlogin connections at the end of an interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/network/tcp/connect.rpt</p>
telnet conn	<p>The number of telnet connections at the end of the interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/network/tcp/connect.rpt</p>
total conn	<p>The total number of connections at the end of the interval [Sequential = AVG Non-Sequential = AVG] View Report: /report/solaris/network/tcp/connect.rpt</p>

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
net interval	The number of seconds elapsed between two data samples of the Network Agent [Sequential = SUM Non-Sequential = ID]
tqbnp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqbnp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

13.4. Workload Statistics

Workload statistics are maintained in the TeamQuest performance database by the Process-Workload Agent (**twarp**). The statistics are classified by the hierarchy of key names.

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
twarp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
twarp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]
warp interval	The number of seconds elapsed between two data samples of the Process-Workload Agent [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

Parameter Hierarchy

Class:	Workload
Subclass:	by Workload
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /workload /workloadset/workload
TeamQuest Table Name:	Workload.by Workload
Open Table Name:	WLBYWORKLOAD
Workload Set:	WLS1, WLS2, ...
Workload:	WL1, WL2, ...
Statistic Name:	
%cpu	<p>The percentage of total CPU consumed by the workload. Total CPU time is sampling interval multiplied by the number of CPUs on the system. Thus, if the sum of %cpu for all workloads is less than 100%, some CPUs must have been idle for some amount of time during the sampling interval.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/pct-cpu.rpt</p>
avgmemb	<p>The cumulative swappable process image size in kilobytes of all of the running processes in the workload at the end of the sampling interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
datafltslp	<p>The data page fault sleep time for the workload. This value only represents the processes active at the end of the sampling interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/sleep.rpt</p>
etime	<p>The sum of the elapsed times in seconds of all of the processes in the workload. Dividing this number by the number of processes in the workload (pongoing + pcomplete) gives the average time a process in the workload existed during the sampling interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
invswch	<p>The number of involuntary context switches for the workload. This value only represents the processes active at the end of the sampling interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
kernfltslp	<p>The kernel page fault sleep time in seconds for the workload. This value only represents the processes active at the end of the sampling interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/sleep.rpt</p>

lioch	<p>The number of logical characters in kilobytes transferred by the workload during the sampling interval [Sequential = SUM Non-Sequential = SUM] View Report: /report/solaris/workload/lioch.rpt</p>
majflt	<p>The number of major page faults generated by the workload for processes that were active at the end of the sampling interval. A major page fault is a page fault that requires I/O. [Sequential = SUM Non-Sequential = SUM] View Report: /report/solaris/workload/maj-flt.rpt</p>
otherslp	<p>The remaining sleep time not accounted for by other sleep time statistics for the workload. This value only represents the processes active at the end of the sampling interval. [Sequential = SUM Non-Sequential = SUM] View Report: /report/solaris/workload/sleep.rpt</p>
pcomplete	<p>The number of processes completed in the sampling interval. For process data, the same number is called cproc. The number of processes in a workload could be derived by adding pongoing and pcomplete, which is represented by the nproc statistic in process data. [Sequential = SUM Non-Sequential = SUM] View Report: /report/solaris/workload/num-proc.rpt</p>
pio	<p>The number of physical I/O transfers done by the workload during the sampling interval [Sequential = SUM Non-Sequential = SUM] View Report: /report/solaris/workload/pio.rpt</p>
pongoing	<p>The number of processes running at the end of the sampling interval. In process data, the same number is called oproc. The number of processes in a workload could be derived by adding pongoing and pcomplete. This sum is called nproc in process data. [Sequential = LST Non-Sequential = SUM] View Report: /report/solaris/workload/num-proc.rpt</p>
prss	<p>The resident set size in kilobytes of private memory occupied by all of the running processes in the workload at the end of the interval [Sequential = AVG Non-Sequential = SUM] View Report: /report/solaris/workload/rss.rpt</p>
pstart	<p>The number of processes started in the sampling interval. In process data, this number is called sproc. [Sequential = SUM Non-Sequential = SUM] View Report: /report/solaris/workload/num-proc.rpt</p>

record_count	<p>The number of collected records represented by the record written to the database. For nonreduced records, this value is 1. For reduced records, this value is the number of records that are combined into a single database record.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
reduction_name	<p>The name of reduction rule</p> <p>[Sequential = ID Non-Sequential = ID]</p>
reduction_source	<p>The source of the reduction record. For reduction records with agent sources, this value is A. For reduction records with harvest sources, this value is H.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
rss	<p>The resident set size (real memory) in kilobytes of all running processes at the end of the interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/rss.rpt</p>
srss	<p>The resident set size in kilobytes of shared memory occupied by all of the running processes in the workload at the end of the interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/rss.rpt</p>
syscpu	<p>The system CPU time in seconds used by the workload. System CPU time is the time spent in kernel mode (for example, the time spent in executing system calls, paging, and so on).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/sys-cpu.rpt</p>
textfltslp	<p>The amount of text page fault sleep time for the workload. This value only represents the processes active at the end of the sampling interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/sleep.rpt</p>
threads	<p>A dynamic object that represents a control point in a process and executes a sequence of instructions</p> <p>[Sequential = LST Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/threads.rpt</p>
totcpu	<p>The total CPU time in seconds used by the workload. This value is the same as the sum of usrcpu + syscpu.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/cpu-util.rpt</p>

userlckslp	<p>The amount of user lock sleep time for the workload. This value only represents the processes active at the end of the sampling interval.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/sleep.rpt</p>
usrcpu	<p>The user CPU time in seconds used by the workload. User CPU time is the time the CPU spent running in user mode.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/user-cpu.rpt</p>
waitcpu	<p>The wait CPU time in seconds used by the workload. Wait CPU time is the time spent waiting for the CPU.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p> <p>View Report: /report/solaris/workload/wait-cpu.rpt</p>

13.5. Process Statistics

The Process-Workload Agent (**tqwarp**) collects process data from the operating system and processes accounting files. The Process-Workload Agent calculates the usage of every process in a given interval, applies reduction definitions to each process, and stores the reduced process data. It also applies workload definitions to the reduced process data and stores system resource usage by workload.

Reduced Process Records

The Process-Workload Agent collects process data and reduces the data according to the user-defined reduction definitions. A reduction definition may cause multiple processes to be merged into a single process record. Thus, a process record contains data about one or more processes. When you are looking at the resource usage numbers, it is important to know how many processes a process record actually represents. When a process record is representing more than one process, the resource usage fields such as `totcpu`, `rss`, and `pio_t` are the sum of the resource usage of the individual processes. When all of the processes do not have the same value for a field, the identifier fields such as `command`, `login`, and `gid` are set to `<Multi>`. When data for some fields is not available, the fields are set to `<N/A>`.

Disabling Reduction Definitions

If you want to look at the details of every individual process and do not wish to have merged process records, you must disable reduction processing by making all reduction sets inactive. However, with reduction processing disabled, more records have to be stored and more disk space is needed. For information on disabling reductions, see the *TeamQuest Performance Software Administration Guide*.

Process Data with Process Accounting Turned Off

If process accounting is turned off, the process data is incomplete as data about completed processes is not available. In this case, the process data only shows a portion of the activity in the sampling interval. To find out whether process accounting is turned on or off, look at the `cproc` field of all of the process records. If the `cproc` field indicates all zeros, it means that processes were not completed in that sampling interval and that the process accounting is turned off.

When process accounting is not turned on, the process record `<Other>` includes the CPU time for processes that completed during the interval.

For more information on process accounting, see the *TeamQuest Performance Software Administration Guide*.

Retrieving Hardware Configuration Information

The Process-Workload Agent retrieves hardware configuration information. The information is stored upon startup and once-a-day in the `HINV.Summary`, `HINV.CPUModel`, `HINV.Devices`, `HINV.FileSystem`, `HINV.FileSystemToDevice`, and `SOLARIS.ZFS Storage Pool` table files of the TeamQuest performance database. The information is also stored if the agent detects a change in configuration.

Table Field Hierarchy

Class:	SOLARIS
Subclass:	Process
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /TeamQuest/System/ <i>systemname</i> /Process
TeamQuest Table Name:	SOLARIS.Process
Open Table Name:	SOLPROC
Collection interval:	Based on the primary aggregation set
Default retention:	1 day
Table type:	Performance

Note: *The collection interval depends on the Processes Only setting in the configuration file for the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.*

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avgmemb_t	The size of the swappable process image in kilobytes. If a process ends within the same interval, the number is unavailable and is marked as <N/A>. [Sequential = AVG Non-Sequential = SUM]
btime	The start time of the process. For process records representing more than one process, this field shows the earliest of the start times. [Sequential = FST Non-Sequential = SUM]
command	The command name of the process. If a process starts and ends within the same interval, only up to 8 characters can be displayed. Otherwise, up to 16 characters are displayed. Therefore, an “automountd” process may appear as “automoun” if it starts and ends within the same interval. In such cases, you may want to set up your workload, reduction, or filter definitions to catch both the command names as in the following: command = {“automoun”, “automountd”} [Sequential = ID Non-Sequential = ID]
cproc	The number of processes completed in the interval [Sequential = SUM Non-Sequential = SUM]
datafltslp	The data page fault sleep time for the current interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
datafltslp_t	The data page fault sleep time since the process started [Sequential = LST Non-Sequential = SUM]

etime	The elapsed time in seconds for the current interval. This number tells how long a process existed in the current interval. [Sequential = SUM Non-Sequential = SUM]
etime_t	The total elapsed time in seconds. This number tells how long a process existed since it started. For a single process, this is the same as the value of the Time field minus the value of the btime field. [Sequential = LST Non-Sequential = SUM]
fullcmd	The full command string, including arguments, for the process. If a process begins and ends in the sampling interval, the field is unavailable and is marked as <N/A>. This is an important consideration when using a workload, reduction, or filter definition. The number of characters that are stored is determined by the Command Length setting of the Process-Workload Agent configuration file in TeamQuest Manager. You can also have either the first or the last <i>N</i> characters of the command displayed. The setting is specified by the Command Orientation setting of the Process-Workload Agent configuration file in TeamQuest Manager. The limit of characters read from the operating system data source is 8095. For a description of how this may affect collection when the LAST Command Orientation is configured for the Process-Workload Agent, see the <i>TeamQuest Performance Software Administration Guide</i> . [Sequential = ID Non-Sequential = ID]
gid	The real group identifier of the owner of the process [Sequential = ID Non-Sequential = ID]
group	The group name of the owner of the process. This field is derived from gid. [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
invswch	The number of involuntary context switches for the current interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
invswch_t	The number of involuntary context switches since the process started [Sequential = LST Non-Sequential = SUM]
kernfltslp	The kernel page fault sleep time for the current interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
kernfltslp_t	The kernel page fault sleep time since the process started [Sequential = LST Non-Sequential = SUM]
lioch	The number of logical characters in kilobytes transferred in the current interval [Sequential = SUM Non-Sequential = SUM]
lioch_t	The total number of logical characters in kilobytes transferred since the start of the process [Sequential = LST Non-Sequential = SUM]

login	The login name of the owner of the process. This field is derived from uid. [Sequential = ID Non-Sequential = ID]
majflt	The number of major page faults generated in the current interval. A major page fault is a page fault that requires I/O. If a process ends in the sampling interval, the number is unavailable and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
majflt_t	The total number of major page faults generated since the process started. A major page fault is a page fault that requires I/O. [Sequential = LST Non-Sequential = SUM]
nproc	The number of processes that the process record represents. If a process starts and ends within the same interval, the number is not available and is marked as <N/A>. In a reduced record, it is the number of processes that were merged together to form a single process record. When time consolidation is not applied to the process data, this number should equal the sum of cproc and oproc. When consolidating over time, the number represents the average number of processes in the process record for the consolidated period. An average is used because processes can move in and out and between reduced records from interval to interval. [Sequential = AVG Non-Sequential = SUM]
oproc	The number of ongoing processes at the end of the interval [Sequential = LST Non-Sequential = SUM]
otherslp	The remaining sleep time not accounted for by other sleep time statistics for the interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
otherslp_t	The remaining sleep time not accounted for by other sleep time statistics since the start of the process. [Sequential = LST Non-Sequential = SUM]
pctcpu	The percentage of total available CPU time the process used in the current sampling interval [Sequential = AVG Non-Sequential = SUM]
pid	The process identifier number. If a process starts and ends within an interval, this number is unavailable and is marked as <N/A>. [Sequential = ID Non-Sequential = ID]
pio	The number of physical I/O requests for the current interval [Sequential = SUM Non-Sequential = SUM]
pio_t	The total number of physical I/O requests since the process started [Sequential = LST Non-Sequential = SUM]
pool	The name of the resource pool that the process is using. If the resource pool is not specified, the value is unavailable and is marked as <N/A>. [Sequential = ID Non-Sequential = ID]
poolid	The resource pool identifier. The default resource pool has a poolid of 0. [Sequential = ID Non-Sequential = ID]

ppid	<p>The numerical identifier of the parent process. If a process starts and ends within an interval, this number is unavailable and is marked as <N/A>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
pri	<p>The priority of the process. Higher numbers mean lower priority. If a process starts and ends within an interval, this number is unavailable and is marked as <N/A>. If the process record represents more than one process, the priority is an average of the constituent running processes' priority values.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
projid	<p>The Solaris Resource Manager project identifier. If extended process accounting is not being used, this value is set to <N/A> for processes that complete in the sampling interval.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
prss	<p>The resident set size of private memory for the process at the end of the interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
pset	<p>The name of the processor set that the process is using. If the processor set is not enabled, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
psetid	<p>The processor set identifier. Solaris assigns a value of -1 to the default processor set.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
redname	<p>The reduction name of the process record. If a process did not match any of the reduction definitions, then it would not be reduced and will not have a reduction name.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
rss	<p>The resident set size (real memory size) of the process. If a process starts and ends within an interval, the number is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
sproc	<p>The number of processes started in the interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
srss	<p>The resident set size of shared memory for the process at the end of the interval. If a process ends within the interval, the value is unavailable and is marked as <N/A>.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
syscpu	<p>The system CPU time in seconds for the current interval. System CPU time is the time the CPU spent running in kernel mode (for example, the time spent in executing system calls, paging, and so on). If an application is taking a lot of syscpu time, you may want to optimize the use of system calls (for example, use a larger block size for I/O).</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>

syscpu_t	The total system CPU time in seconds [Sequential = LST Non-Sequential = SUM]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
textfltslp	The amount of text page fault sleep time for the current interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
textfltslp_t	The amount of text page fault sleep time since the process started [Sequential = LST Non-Sequential = SUM]
threads	The number of threads associated with the process at the end of the interval. This value represents the number of light-weight processes (LWP) or kernel-supported user threads. A thread is a dynamic object that represents a control point in a process and executes a sequence of instructions. [Sequential = LST Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totcpu	The total CPU time in seconds used in the current interval. This number is the same as the sum of usrcpu and syscpu. [Sequential = SUM Non-Sequential = SUM]
totcpu_t	The total CPU time (user + system) in seconds used by the process since it started. This number is the same as the sum of usrcpu_t and syscpu_t. [Sequential = LST Non-Sequential = SUM]
tty	The controlling terminal identifier in dev_t format. For the processes without a controlling terminal, this field will contain a -1. [Sequential = ID Non-Sequential = ID]
ttyname	The controlling terminal for the process. It is a device name without the /dev/ prefix. This is derived from tty. For the processes without a controlling terminal, this field will contain a question mark (?). [Sequential = ID Non-Sequential = ID]
uid	The real user id of the process owner [Sequential = ID Non-Sequential = ID]
userlckslp	The amount of user lock sleep time for the current interval. If a process ends within the sampling interval, the value is not available and is marked as <N/A>. [Sequential = SUM Non-Sequential = SUM]
userlckslp_t	The amount of user lock sleep time since the process started [Sequential = LST Non-Sequential = SUM]
usrcpu	The user CPU time in seconds for the current interval. User CPU time is the time the CPU spent running in user mode. If an application is taking a lot of usrcpu time, you should try to optimize the code, if possible. [Sequential = SUM Non-Sequential = SUM]

usrcpu_t	The total user CPU time in seconds since the start of the process [Sequential = LST Non-Sequential = SUM]
waitcpu	The wait CPU time in seconds for the current interval. Wait CPU time is the time spent waiting for the CPU. [Sequential = SUM Non-Sequential = SUM]
waitcpu_t	The total wait CPU time in seconds since the start of the process [Sequential = LST Non-Sequential = SUM]
Workload	<p>The workload set and the workload associated with the process. When the Process-Workload Agent stores the process record, this field is blank. When process records are reported, the workload can be evaluated and is shown in the report.</p> <p>This field is available for reporting only when using TeamQuest Analyzer and TeamQuest tView.</p> <p>Workload evaluation takes place when data is retrieved from the database, based on workload sets defined in the database where the data is stored. Workload sets reported in the Workload field do not have to be active.</p> <p>For more information on workload evaluation, see the <i>TeamQuest Analyzer User Guide</i> or the <i>TeamQuest Performance Software Command Line Interfaces Reference Manual</i>. [Sequential = ID Non-Sequential = ID]</p>
workload:wlsname	<p>There is one field for each <i>wlsname</i> (Workload Set Name). The value for this field shows the name of the workload to which the process belongs. If a process belongs to none of the workloads, it will display the workload name "OTHER."</p> <p>This field is available for reporting only when using TeamQuest View or TeamQuest cView. [Sequential = ID Non-Sequential = ID]</p>
zone	<p>The zone name of the zone where the process is running. If extended process accounting is not being used, this value is set to <N/A> for processes that complete in the sampling interval. [Sequential = ID Non-Sequential = ID]</p>

13.6. Project Statistics

Project statistics are stored in the TeamQuest performance database tables. Project statistics are available when a Solaris project has been modified.

Table Field Hierarchy

Class:	SOLARIS
Subclass:	Project
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	SOLARIS.Project
Open Table Name:	SOLPROJECT
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
attributes	The semicolon-separated list of name value pairs. Each pair is in the form <i>name</i> [= <i>value</i>], where <i>name</i> is the key name and <i>value</i> is the optional key value. [Non-Sequential = ID]
comment	The project description. This field is blank if there is no project description. [Non-Sequential = ID]
cpushares	The number of project cpu shares for the Fair Share Scheduler (FSS) [Non-Sequential = SUM]
groups	The comma-separated list of groups of users allowed in the project. A blank field indicates that no groups are allowed in the project. [Non-Sequential = ID]
pool	The name of the resource pool to which the process is assigned. If no resource pool is assigned, the value is not available and is marked as <N/A>. [Non-Sequential = ID]
project	The name of the project [Non-Sequential = ID]
projid	The unique numerical identifier of the project. Project IDs below 100 are reserved for the operating system. [Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Non-Sequential = ID]

Time	The timestamp of the data sample [Non-Sequential = ID]
users	The comma-separated list of users allowed in the project. This field is blank if no users are allowed in the project. [Non-Sequential = ID]

13.7. Hardware Inventory Statistics

The Process-Workload Agent (**tqwarp**) retrieves hardware configuration information. The information is stored upon startup and once-a-day in the HINV.Summary, HINV.CPUModel, HINV.CPU Thread Speeds, HINV.Devices, HINV.FileSystem, HINV.FileSystemToDevice, and SOLARIS.ZFS Storage Pool table files of the TeamQuest performance database. The information is also stored if the agent detects a change in configuration.

Note: *The storage of hardware inventory records is also dependent on the Hardware Inventory setting in the configuration file of the Process-Workload Agent. For more information, see the section on configuring the Process-Workload Agent in the TeamQuest Performance Software Administration Guide.*

Table Field Hierarchy

Class:	HINV
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	HINV.Summary
Open Table Name:	HINVSUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
core_multi_thread	The status or ability of the processor to support multiple independent threads. The field will contain <N/A> if the information is not available. [Non-Sequential = ID]
cores_per_chip	The number of cores or processors on an individual chip. The value will be zero if the information is not available. [Non-Sequential = ID]
cpu_chips	The number of CPU chips or sockets. The value will be zero if the information is not available. [Non-Sequential = ID]
cpu_count	The number of configured processors [Non-Sequential = ID]
cpu_speed	The speed of the processor in MHz or GHz [Non-Sequential = ID]

cpu_type	The basic instruction set architecture of the current system [Non-Sequential = ID]
logical_cpu_count	The number of logical processors [Non-Sequential = ID]
mem_size	The size of configured random access memory in kilobytes, where 1 kilobyte = 1,024 bytes [Non-Sequential = ID]
memory	The size of configured random access memory in megabytes, where 1 megabyte = 1,048,576 bytes [Non-Sequential = ID]
memory_size	The size of configured random access memory in megabytes or gigabytes [Non-Sequential = ID]
model	The name of the hardware implementation or platform [Non-Sequential = ID]
os_release	The name and level of the implementation of the operating system [Non-Sequential = ID]
pagesize	The size of a page of memory [Non-Sequential = ID]
partition_type	The partition type of the system. The value indicates the system hypervisor type, guest type, logical partition type, zone type, or logical domain type. If the system does not have a partition type, this field will be blank. [Non-Sequential = ID]
serial	The hardware-specific serial number of the physical machine [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_identifier	The information used to identify the system [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
timezone	The time zone where the data was collected [Non-Sequential = ID]
TQLevel	The level of TeamQuest Manager [Non-Sequential = ID]

The HINV.CPUModel table stores best-match, relative performance data about the system configuration. This table is created by the Hardware Inventory Agent (**tqhin**) to map physical hardware to a CPU model that describes performance in relative terms. This table is not created for any virtualized system. It is populated for physical systems only. It is not populated for VMware guests, Hyper-V guests, Solaris logical domains (LDOMs), Solaris guest LDOMs, KVM guests, and Linux on POWER systems.

Table Field Hierarchy

Class:	HINV
Subclass:	CPUModel
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	HINV.CPUModel
Open Table Name:	HINVCPUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
cpu_chips	The number of CPU chips or sockets [Non-Sequential = ID]
cpu_confidence	The percentage of confidence in the correctness of the CPU match based on model, frequency, and configuration (chips, cores, threads) [Non-Sequential = SUM]
cpu_cores	The number of CPU cores or processors on an individual CPU chip [Non-Sequential = ID]
cpu_name	The name of the selected CPU [Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Non-Sequential = ID]
cpu_speed	The speed of the processor in megahertz (MHz) or gigahertz (GHz) [Non-Sequential = ID]
cpu_threads	The number of CPU threads on an individual CPU core or processor [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
user_override	The user override status of the default TeamQuest generated CPU match. This field is not currently used and should appear as 0. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	CPU Thread Speeds
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.CPU Thread Speeds
Open Table Name:	HINVCPUTHREADSPEEDS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
speed_up_factor	The performance improvement when there are multiple active threads per core, compared to when there is only one active thread per core [Sequential = AVG Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
thread_number	The number of active threads [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = ID Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	Devices
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.Devices
Open Table Name:	HINVDEVS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
class	The device classification: controller, disk, or tape [Non-Sequential = ID]
controller	The device path indicator which defines a connection to another device [Non-Sequential = ID]
lun_id	The globally unique Logical Unit Number (LUN) identifier for Storage Area Network (SAN) based disk devices. This field is blank for non-SAN based disk devices, CD-ROM drives, tape drives, and so on. [Non-Sequential = ID]
name	The unique identifier for the device [Non-Sequential = ID]
name2	The alternate device name. This field may be blank. [Non-Sequential = ID]
product	The product identifier. This field may be blank. [Non-Sequential = ID]
revision	The revision level for the product. This field may be blank. [Non-Sequential = ID]
rpm	The speed at which the media spins. If an actual value cannot be obtained for the device, a default value of 7,200 is used. [Non-Sequential = ID]
sequence	The sequence number of the device [Non-Sequential = ID]
swap	A true or false statement which indicates whether or not a swap file exists on the device [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]

Time	The timestamp of the data sample [Non-Sequential = ID]
vendor	The name of the device vendor. This field may be blank. [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	FileSystem
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.FileSystem
Open Table Name:	HINVFILESYS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
BlkSize	The size of a block on the file system [Non-Sequential = ID]
Device	The path for the device on which the file system is mounted [Non-Sequential = ID]
Name	The unique identifier for the file system [Non-Sequential = ID]
Source	The source logical volume of the file system. For Zettabyte File System (ZFS), this field contains the name of the ZFS storage pool that hosts the file system. This field will be blank for file systems that use a physical disk or physical disk partition directly. [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
TotBlks	The total number of blocks on the file system [Non-Sequential = ID]
TotFiles	The maximum total number of files, as represented by inodes, possible on the file system. Some inodes may be used for entities other than visible files. [Non-Sequential = ID]
TotSize	The total amount of space on the file system in megabytes [Non-Sequential = ID]
Type	The type of the file system [Non-Sequential = ID]

Table Field Hierarchy

Class:	HINV
Subclass:	FileSystemToDevice
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.FileSystemToDevice
Open Table Name:	HINVFSTODEV
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
Device	The name of the device [Non-Sequential = ID]
FileSystem	The name of the file system [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

Table Field Hierarchy

Class:	SOLARIS
Subclass:	ZFS Storage Pool
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	SOLARIS.ZFS Storage Pool
Open Table Name:	SOLZFSPOOL
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
allocated	The total size of all allocated file systems on the Zettabyte File System (ZFS) pool [Non-Sequential = SUM]
altroot	The alternative root directory of the ZFS pool. If set, this directory is prepended to any mount points within in the ZFS pool. [Non-Sequential = ID]

autoexpand	The automatic expand setting. If set to on, the ZFS pool automatically grows in size when the underlying Logical Unit Number (LUN) backing the pool grows in size. [Non-Sequential = ID]
autoreplace	The automatic replacement setting. If set to on, any new device found in the same physical location as a device that previously belonged to the pool is automatically formatted and added to the ZFS pool. [Non-Sequential = ID]
bootfs	The default bootable dataset for the root ZFS pool [Non-Sequential = ID]
cachefile	The location of the pool configuration cache file. This file is used to discover ZFS pools at system boot time. [Non-Sequential = ID]
capacity	The percentage of the ZFS pool that is allocated (used) [Non-Sequential = AVG]
dedupditto	The threshold for making multiple copies of the duplicated data blocks. If the reference count for a deduplicated block goes above this threshold, another copy of the block is stored automatically. [Non-Sequential = ID]
dedupratio	The deduplication ratio specified for the ZFS pool. This value is expressed as a multiplier. For example, a dedupratio value of 1.3 indicates that 1.3 units of data were stored but only 1 unit of disk space was actually consumed. [Non-Sequential = ID]
delegation	The user access setting for accessing the ZFS pool. If set to on, non-privileged users are granted access to the ZFS pool based on dataset permissions. [Non-Sequential = ID]
failmode	The action to be taken if the ZFS pool has a catastrophic failure. Values for this field include: wait, continue, and panic. [Non-Sequential = ID]
free	The amount of free space remaining for allocation in the ZFS pool [Non-Sequential = SUM]
guid	The globally unique identifier of the ZFS pool [Non-Sequential = ID]
health	The health state of the ZFS pool. Values for this field include: DEGRADED, FAULTED, OFFLINE, ONLINE, REMOVED, or UNAVAIL. [Non-Sequential = ID]
listsnapshots	The list of snapshots associated with the ZFS pool. If set to on, information about the ZFS pool snapshots are displayed. Snapshots can be listed in the “zfs list” command. [Non-Sequential = ID]
name	The name of the ZFS pool [Non-Sequential = ID]

readonly	The read only status of the ZFS pool. If set to on, the ZFS pool operates in a read only mode and the configuration of the ZFS pool cannot be changed. Similarly, any mounted file systems associated with the ZFS pool operate in read only mode. [Non-Sequential = ID]
size	The total size of the ZFS pool [Non-Sequential = SUM]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
version	The current on-disk version of the ZFS pool [Non-Sequential = ID]

13.8. System Log Statistics

The System Log Agent (**tqslp**) is used to collect system log messages generated by the system log daemon (syslogd). The System Log Agent stores these messages in the TeamQuest performance database for analysis and alarm reporting. The log messages are separated into four fields; the time that the message was posted, the host system from where the message was initiated, the program or user that posted the message, and the text of the message.

Table Field Hierarchy

Class:	System
Subclass:	System Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /System Log
TeamQuest Table Name:	System.System Log
Open Table Name:	SYSSYSTEMLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Event_Time	The time that the message was logged to the system log [Non-Sequential = ID]
Loghost	The name of the system that logged the message [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Reporter	The name of the user or process that logged the message [Non-Sequential = ID]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system where the log message originated. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

13.9. General Log Statistics

The General Log Agent (**tqglp**) is used to collect log messages generated by application programs. The General Log Agent stores these messages in the TeamQuest performance database for analysis. Examples include backup, security, database, and Web server applications.

Table Field Hierarchy

Class:	System
Subclass:	General Log
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /General Log
TeamQuest Table Name:	System.General Log
Open Table Name:	SYSGENERALLOG
Collection interval:	N/A
Default retention:	4 days
Table type:	Event

Statistic Name	Description
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The message type [Non-Sequential = ID]

13.10. TeamQuest Log Statistics

The following statistics are stored in the performance database tables by the TeamQuest Log Agent (**tqlog**). The collection interval and retention periods can be modified. For more information on modifying the collection interval and retention periods, see the *TeamQuest Performance Software Administration Guide*.

Table Field Hierarchy

Class:	Service
Subclass:	TeamQuest Log
IT Resource Name:	/TeamQuest/System/systemname/TeamQuest Log
TeamQuest Table Name:	Service.TeamQuest Log
Open Table Name:	SVCTQLOG
Collection interval:	N/A
Default retention:	1 day
Table type:	Event

Statistic Name	Description
Filename	The name of the TeamQuest log file that was the source of the message text [Non-Sequential = ID]
Message	The message text [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Sequence	The sequence number of the message in the sampling interval [Non-Sequential = ID]
System	The name of the system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	The log message type. This is always set to tqlog . [Non-Sequential = ID]

13.11. Derived Statistics

Some products within TeamQuest Performance Software use derived statistics to display common statistics across different platforms. The derived statistics are inserted into the performance database when databases are created. In this section, a derived statistic is marked with an asterisk (*).

See the following statistics for more information:

- Workload Performance Derived Statistics (see 13.11.1)
- TeamQuest On the Web Derived Statistics (see 13.11.2)
- TeamQuest Alert Derived Statistics (see 13.11.3)
- Rules Derived Statistics (see 13.11.4)

13.11.1. Workload Performance Derived Statistics

TeamQuest Manager maintains derived statistics that use data from the System Activity Agent (**tqbsp**) and Process-Workload Agent (**tqwarp**). The workload performance reports reference these statistics. For information on workload performance reports, see the *TeamQuest View Reports Reference Manual*.

Parameter Hierarchy

Class:	Derived
Subclass:	Workload Performance.by Workload
Workload Set:	WLS1, WLS2, ...
Workload:	ALL
Statistic Name:	
%cpu*	The total percentage of CPU utilization. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt
Kbytes resident memory/process*	The average amount of resident memory used per process. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt
Kbytes virtual memory/process*	The average amount of virtual memory used per process. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt
PIOs/sec*	The number of physical I/Os per second. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt

Population (etime/interval)*	The average number of concurrent processes. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt
Response (etime/process)*	The elapsed time per process. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt
Throughput (processes/sec)*	The number of processes completed per second. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt
Total Kbytes resident memory*	The average amount of resident memory used. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt
Total Kbytes virtual memory*	The average amount of virtual memory used. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/workload.rpt
Class:	Derived
Subclass:	Workload Performance.Summary
Workload Set:	WLS1, WLS2, ...
Statistic Name:	
%cpu*	The total percentage of CPU utilization. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt
block IO avresp*	The average disk and tape I/O response in milliseconds. Collected by the System Activity Agent. View Report: /report/solaris/wkldperf/overall.rpt
block IO r+w/s*	The number of disk and tape I/Os per second. Collected by the System Activity Agent. View Report: /report/solaris/wkldperf/overall.rpt
Kbytes resident memory/process*	The average amount of resident memory used per process. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt
Kbytes virtual memory/process*	The average amount of virtual memory used per process. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt

PIOs/sec*	The number of physical I/Os per second. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt
Population (etime/interval)*	The average number of concurrent processes. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt
Response (etime/process)*	The elapsed time per process. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt
Throughput (processes/sec)*	The number of processes completed per second. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt
Total Kbytes resident memory*	The average amount of resident memory used. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt
Total Kbytes virtual memory*	The average amount of virtual memory used. Collected by the Process-Workload Agent. View Report: /report/solaris/wkldperf/overall.rpt

13.11.2. TeamQuest On the Web Derived Statistics

The derived statistics used by TeamQuest On the Web include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQWeb.Summary
Statistic Name:	
avg_disk_queue_length*	The average number of requests outstanding for all of the devices
avg_service_time*	The average time in milliseconds for a transfer request to be completed
buffer_pct_read_cache*	The percentage of logical reads satisfied from the buffer cache
buffer_pct_write_cache*	The percentage of logical writes satisfied from the buffer cache
disk_xfers_per_sec*	The total number of read and write transfers per second for all of the devices
free_disk_space*	The amount of space available (not in use) on all file systems in megabytes. This measurement is taken at the end of the sampling interval and includes the space held back from normal users.
free_swap_space*	The number of megabytes free for process swapping

<code>free_real_mem*</code>	The amount of free memory available in megabytes. This measurement is taken at the end of the sampling interval.
<code>nfs_calls_per_sec*</code>	The total number of NFS calls sent by the client
<code>page_ins_per_sec*</code>	The number of swap-in requests per second
<code>page_outs_per_sec*</code>	The number of page-out requests per second
<code>page_scans_per_sec*</code>	The number of pages per second the page daemon scans to see if they can be freed
<code>pct_cpu_busy*</code>	The percentage of total CPU time that the CPU was not idle. This value includes the time running system code and the time running normal priority user processes.
<code>pct_disk_busy*</code>	The percentage of time a disk was busy servicing a transfer request
<code>pct_sys_cpu*</code>	The percentage of total CPU time spent in system mode
<code>pct_usr_cpu*</code>	The percentage of total CPU time spent running in user mode
<code>pkt_errors_per_sec*</code>	The total number (in + out) of network errors per second for all network interfaces
<code>pkts_in_per_sec*</code>	The total number of network input packets per second for all network interfaces
<code>pkts_out_per_sec*</code>	The total number of network output packets per second for all network interfaces
<code>pkts_per_sec*</code>	The total number (in + out) of network packets per second for all network interfaces
<code>total_disk_space*</code>	The total (used + available) amount of space on all file systems in megabytes. This measurement is taken at the end of the sampling interval and includes the space held back from normal users.
<code>total_processes*</code>	The number of entries currently being used in the process table. This measurement is taken at the end of the sampling interval.
<code>total_real_mem*</code>	The total amount of real (physical) memory in megabytes. This measurement is taken at the end of the sampling interval.
<code>total_swap_space*</code>	The total number of megabytes available for swapping

13.11.3. TeamQuest Alert Derived Statistics

The derived statistics used by TeamQuest Alert include the following:

Parameter Hierarchy

Class:	Derived
Subclass:	TQAlert.Summary
Statistic Name:	
free_real_mem*	The average amount of memory available to user processes in megabytes
kmem_fails*	The number of kernel requests of memory that were not satisfied
net_errors*	The number of network errors for all network interfaces
page_scans*	The number of pages per second scanned by the page-stealing daemon
pct_swap_free*	The percentage of unused swap space at the end of the sampling interval in megabytes
pct_wio*	The percentage of time spent idle while processes are waiting for I/O completion
run_queue*	The average length of the run queue (a queue of processes in memory and runnable) while the run queue is occupied
total_processes*	The total number of processes active on the system

13.11.4. Rules Derived Statistics

The TeamQuest View Rules reports reference the following derived statistics. For more information, see the *TeamQuest View Reports Reference Manual*.

Parameter Hierarchy

Class:	Derived
Subclass:	Rules.CPU Rules
Statistic Name:	
mutex_warning*	The CPU mutex rate divided by 500
pct_cpu_busy*	The percentage of time the user and system CPU was not idle View Report: /report/solaris/rules/cpu_usage.rpt
Class:	Derived
Subclass:	Rules.File System Rules
Statistic Name:	
%inodes free*	The percentage of total i-nodes that are free View Report: /report/solaris/rules/inodes_free.rpt
user space free*	The amount of free user disk space View Report: /report/solaris/rules/file_space.rpt
Class:	Derived
Subclass:	Rules.Network Rules
Statistic Name:	
%collisions*	The percentage of network interface out packets that had collisions View Report: /report/solaris/rules/network_collisions.rpt
%defer*	The percentage of network interface out packets that were deferred View Report: /report/solaris/rules/network_defer.rpt

Class: Derived
 Subclass: Rules.NFS Rules

Statistic Name:

maxtimeout* A maximum time out value for client RPC calls
 View Report:
 /report/solaris/rules/nfs.rpt

rpc_client_calls* The total number of client RPC calls
 View Report:
 /report/solaris/rules/nfs.rpt

rpc_client_timeouts* A total number of client RPC timeouts
 View Report:
 /report/solaris/rules/nfs.rpt

Class: Derived
 Subclass: Rules.Ram Rules

Statistic Name:

maxrestime* The maximum amount of idle page residence time. Used to avoid the possibility of dividing by zero when deriving idle page residence time (restime).

minscans* The minimum amount of ipage scan rate. Used to avoid the possibility of dividing by zero when deriving idle page residence time (restime).

restime* An approximation to idle page residence time. The time that an idle page will stay in memory.
 View Report:
 /report/solaris/rules/ram.rpt

13.12. Zone Statistics

Zone statistics are maintained in the TeamQuest performance database by the Process-Workload Agent (**tqwarp**). Zone statistics are updated when a new zoneid appears in the process table or when an existing zoneid no longer appears in the process table.

Table Field Hierarchy

Class:	SOLARIS
Subclass:	Zone
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	SOLARIS.Zone
Open Table Name:	SOLZONE
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
brand	The brand zone of the operating system [Non-Sequential = ID]
capped-cpu.ncpus	The percentage of CPU time that can be used by a zone. The value relates to the zone.cpu-cap resource control. If the CPU time is not specified, the value is not available and is marked as <N/A>. [Non-Sequential = ID]
capped-memory.locked	The maximum amount of locked memory for the zone in megabytes. This value relates to the zone.max-locked-memory resource control. If the amount of locked memory is not specified, the value is not available and is marked as <N/A>. [Non-Sequential = ID]
capped-memory.physical	The maximum amount of physical memory for the zone in megabytes. This value relates to the zone.max-rss resource control. If the amount of physical memory is not specified, the value is unavailable and is marked as <N/A>. [Non-Sequential = ID]
capped-memory.swap	The maximum amount of swap memory for the zone in megabytes. This value relates to the zone-max-swap resource control. If the amount of swap memory is not specified, the value is unavailable and is marked as <N/A>. [Non-Sequential = ID]
cpu-shares	The number of zone cpu shares for the Fair Share Scheduler (FSS). If the number of zone cpu shares is not specified, the value is unavailable and is marked as <N/A>. [Non-Sequential = ID]

dedicated-cpu.importance	The importance property for the temporary resource pool that is created when the zone is booted. If the importance property is not specified, the value is unavailable and is marked as <N/A>. [Non-Sequential = LST]
dedicated-cpu.ncpus	The number of CPUs that should be assigned to the zone. This value can be a range of the minimum and maximum number of processors in the temporary processor set that is created when the zone is booted. If the number of CPUs that should be assigned to the zone is not specified, the value is unavailable and is marked as <N/A>. [Non-Sequential = ID]
id	The zone identifier. A value of <N/A> indicates the zone is not active. [Non-Sequential = ID]
ip	The IP type of the zone. The value can be shared or exclusive. [Non-Sequential = ID]
path	The path to the zone file system [Non-Sequential = ID]
pool	The name of the resource pool that was bound to the zone when the zone was booted. If no resource pool is assigned, the value is not available and is marked as <N/A>. [Non-Sequential = ID]
scheduling-class	The scheduling class that is set when the zone is booted. It is possible to override the process scheduling class for processes running in a zone. The scheduling class can be Fair Share Scheduling (FSS), Time Share (TS), or Real Time (RT). If the scheduling class is not specified, the value is unavailable and is marked as <N/A>. [Non-Sequential = ID]
state	The state of the zone. This value can be configured, incomplete, ready, installed, or running. [Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
uuid	The zone unique universal identifier. This value can be optionally attached to a zone and does not change if the zone is renamed or moved. [Non-Sequential = ID]
zone	The name of the zone. The name of the default zone is global. [Non-Sequential = ID]

13.13. Processor Set Statistics

Processor set statistics are stored in the TeamQuest performance database by the Process-Workload Agent (**tqwarp**). Processor set statistics are available if the resource pool facility is turned on. The statistics are updated upon startup, once every 24 hours, and when a process using a new processor set is detected.

Table Field Hierarchy

Class:	SOLARIS
Subclass:	Processor Set
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	SOLARIS.Processor Set
Open Table Name:	SOLPROCSET
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
comment	The description of the processor set. This field is blank if there is no processor set description. [Non-Sequential = ID]
max	The maximum number of CPUs that can be assigned to the processor set [Non-Sequential = ID]
min	The minimum number of CPUs that can be assigned to the processor set. If this number is more than the number of available processors, the processor set and associated resource pools or zones are not available. If a value of 0 is displayed, processes assigned to the processor set will only get CPU time when the processor is not required by other processor sets. [Non-Sequential = ID]
poold_objectives	The objectives for the pool resource controller (poold) as stated in the configuration of the process set [Non-Sequential = ID]
pset	The name of the processor set [Non-Sequential = ID]
pset_default	Indicates the designated default processor set. If a value of TRUE is displayed, the processor set is the default processor set. If a value of FALSE is displayed, the processor set is not the default processor set. [Non-Sequential = ID]
psetid	The processor set identifier. Solaris assigns a value of -1 to the default processor set. [Non-Sequential = ID]

System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Non-Sequential = ID]
temporary	The status of the processor set. If a value of TRUE is displayed, the processor set is temporary. If a value of FALSE is displayed, the processor set is permanent. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

13.14. Resource Pool Statistics

Resource pool statistics are stored in the TeamQuest performance database by the Process-Workload Agent (**tgwarp**). Resource pool statistics are available if the resource pool facility is turned on. The statistics are updated upon startup, once every 24 hours, and when a process using a new resource pool is detected.

Table Field Hierarchy

Class:	SOLARIS
Subclass:	Resource Pool
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	SOLARIS.Resource Pool
Open Table Name:	SOLRESPOOL
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
comment	The description of the resource pool. This field is blank if there is no resource pool description. [Non-Sequential = ID]
important	The relative importance of the pool. This value can be used for possible resource dispute resolution by the resource pool controller, called poold. [Non-Sequential = ID]
pool	The name of the resource pool [Non-Sequential = ID]
pool_active	The active status of the resource pool. If a value of TRUE is displayed, the resource pool is active. If a value of FALSE is displayed, the resource pool is not active. [Non-Sequential = ID]

pool_default	Indicates the designated default resource pool. If a value of TRUE is displayed, the resource pool is the default resource pool. If a value of FALSE is displayed, the resource pool is not the default resource pool. [Non-Sequential = ID]
poolid	The resource pool identifier. The default poolid is 0. [Non-Sequential = ID]
pset	The name of the processor set with which the resource pool is associated. There can only be one processor set associated with a resource pool, but multiple pools can be associated with a processor set. [Non-Sequential = ID]
scheduler	The specific scheduling class that the resource pool uses when allocating CPU resources to processes working in the project or a zone in which the resource pool is assigned. This field will be blank if the scheduling class is not specified. [Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Non-Sequential = ID]
temporary	The status of the resource pool. If a value of TRUE is displayed, the resource pool is temporary. If a value of FALSE is displayed, the resource pool is permanent. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]

Section 14

PostgreSQL Database Agent

The PostgreSQL Database Agent (**tpggsq1**) collects information about the PostgreSQL database usage and gathers usage statistics on each table in the database.

This section contains a listing of the statistics collected by the agent:

- Database Summary Statistics (see 14.1)
- Database Detail Statistics (see 14.2)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV =Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

14.1. Database Summary Statistics

Summary information about the PostgreSQL database is stored within the PostgreSQL.Database Summary table in the PostgreSQL database.

Table Field Hierarchy

Class:	PostgreSQL
Subclass:	Database Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	PostgreSQL.Database Summary
Open Table Name:	PGSQLDBSUM
Collection interval:	5 minutes (default)
Default retention:	N/A
Table type:	Performance

Statistic Name	Description
active_conn	The number of active connections to the database [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
database	The name of the database. This field is limited to 63 characters. Any database name longer than 63 characters will be truncated. [Sequential = ID Non-Sequential = ID]
free_disk_mb	The amount of free disk space in megabytes (MB) on the file system where the database is located [Sequential = AVG Non-Sequential = SUM]
index_size_mb	The size of all the indexes in megabytes (MB) for the database [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
pct_free_disk	The parentage of free space on the file system where the database is located [Sequential = AVG Non-Sequential = SUM]
System	The name of the system where the agent is running. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
table_size_mb	The size of all the tables in megabytes (MB) in the database [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

total_disk_mb	The total amount of disk space in megabytes (MB) on the file system where the database is located [Sequential = AVG Non-Sequential = SUM]
total_size_mb	The size of all the indexes and tables in megabytes (MB) in the database [Sequential = AVG Non-Sequential = SUM]

14.2. Database Detail Statistics

Detailed information about the PostgreSQL database is stored within the PostgreSQL.Table Detail table in the PostgreSQL database.

Table Field Hierarchy

Class:	PostgreSQL
Subclass:	Table Detail
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	PostgreSQL.Table Detail
Open Table Name:	PGSQLTBLDETAIL
Collection interval:	1 hour (default)
Default retention:	N/A
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
database	The name of the database. This field is limited to 63 characters. Any database name longer than 63 characters will be truncated. [Sequential = ID Non-Sequential = ID]
dead_tuples	The number of dead tuples for the table. A tuple is a record in the database. [Sequential = AVG Non-Sequential = SUM]
index_ratio	The index ratio for the table. Calculated as $\text{index ratio} = \text{index_size_mb} / \text{table_size_mb}$ [Sequential = AVG Non-Sequential = AVG]
index_size_mb	The size of all the indexes in megabytes (MB) for the table [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
last_analyzed	The date the table was last analyzed [Sequential = LST Non-Sequential = ID]

last_autoanalyzed	The date the table was last automatically analyzed [Sequential = LST Non-Sequential = ID]
last_autovacuumed	The date the table was last automatically vacuumed [Sequential = LST Non-Sequential = ID]
last_vacuumed	The date the table was last vacuumed [Sequential = LST Non-Sequential = ID]
live_tuples	The number of live tuples for the table. A tuple is a record in the database. [Sequential = AVG Non-Sequential = SUM]
schema_name	The name of the schema for the table [Sequential = ID Non-Sequential = ID]
System	The name of the system where the agent is running. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
table_name	The name of the table [Sequential = ID Non-Sequential = ID]
table_size_mb	The size of the tables in the database [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
total_size_mb	The total size of the table and all the indexes in megabytes (MB) for the table [Sequential = AVG Non-Sequential = SUM]

Section 15

Sybase ASE Server

Note: *These statistics are only available for use with Sybase ASE level 15.*

The Sybase ASE Agent (**tqsybase**) collects performance data from Sybase servers. The agent obtains parameter data pertaining to CPU, disk capacity, disk I/O, locks, memory, networks, procedure cache, SQL cache, system, and transaction summaries. It also gathers table data pertaining to active SQL, configuration, database detail and summary, device detail, engine detail, lock detail, process, process waits, and system waits statistics.

This section contains a listing of the statistics collected by the agent:

- Performance Statistics (see 15.1)
- Active SQL Statistics (see 15.2)
- Configuration Statistics (see 15.3)
- Database Detail Statistics (see 15.4)
- Database Summary Statistics (see 15.5)
- Device Detail Statistics (see 15.6)
- Engine Detail Statistics (see 15.7)
- Lock Detail Statistics (see 15.8)
- Process Statistics (see 15.9)
- Process Waits Statistics (see 15.10)
- System Waits Statistics (see 15.11)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

15.1. Performance Statistics

All aggregation set parameters for Sybase ASE servers will be classified within the Sybase ASE Category Group key. The actual system name from which the data is collected will be used in the System key. The remaining keys are defined in the following tables:

Parameter Hierarchy

Class:	Sybase ASE
Subclass:	CPU
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.CPU
Open Table Name:	SYBASECPU
Resource:	instance1, instance2, ...
Statistic Name:	
%CpuBusy	The percentage of time that the Adaptive Server is using the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/cputil.rpt
%Idle	The percentage of time that the Adaptive Server is not using the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/cputil.rpt
%IoBusy	The percentage of time that the Adaptive Server is doing I/O operations [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/cputil.rpt
CpuBusy	The total amount of CPU in seconds that the Adaptive Server was performing Adaptive Server work [Sequential = SUM Non-Sequential = SUM]
Idle	The total amount of CPU in seconds that the Adaptive Server is idle [Sequential = SUM Non-Sequential = SUM]
IoBusy	The total amount of CPU in seconds that the Adaptive Server is performing I/O operations [Sequential = SUM Non-Sequential = SUM]
TotalCpuTics	The sum of CPU ticks used for all engines [Sequential = SUM Non-Sequential = SUM]

Sybase ASE Server

Class:	Sybase ASE
Subclass:	Disk.Capacity
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Disk.Capacity
Open Table Name:	SYBASEDISKCAPACITY
Resource:	instance1, instance2, ...
Statistic Name:	
%FreeDisk	The percentage of disk space that is not being used [Sequential = LST Non-Sequential = SUM]
%UsedDisk	The percentage of disk space that is being used [Sequential = LST Non-Sequential = SUM]
FreeDiskMB	The amount of allocated disk space in megabytes (MB) that is not being used [Sequential = LST Non-Sequential = SUM] View Report: /report/sybaseASE/diskutil.rpt
TotalDiskMB	The amount of allocated disk space in megabytes (MB) [Sequential = LST Non-Sequential = SUM] View Report: /report/sybaseASE/diskutil.rpt
UsedDiskMB	The amount of allocated disk space in megabytes (MB) that is being used [Sequential = LST Non-Sequential = SUM]
Class:	Sybase ASE
Subclass:	Disk.IO
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Disk.IO
Open Table Name:	SYBASEDISKIO
Resource:	instance1, instance2, ...
Statistic Name:	
DiskErrors/s	The number of Adaptive Server errors per second when reading and writing to disk [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/ioutil.rpt
DiskReads/s	The number of Adaptive Server reads to disk per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/ioutil.rpt
DiskWrites/s	The number of Adaptive Server writes to disk per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/ioutil.rpt

Class:	Sybase ASE
Subclass:	Locks
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Locks
Open Table Name:	SYBASELOCKS
Resource:	instance1, instance2, ...
Statistic Name:	
%Deadlocks	The percentage of deadlocks as a percentage of the total number of locks [Sequential = AVG Non-Sequential = AVG]
AvgLockContention	The average number of times there was lock contention as a percentage of the total number of lock requests [Sequential = AVG Non-Sequential = AVG]
LockRequests/s	The number of lock requests per second [Sequential = SUM Non-Sequential = SUM]
TotalLockRequests	The total number of lock requests in the interval [Sequential = SUM Non-Sequential = SUM]
Class:	Sybase ASE
Subclass:	Memory
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Memory
Open Table Name:	SYBASEMEM
Resource:	instance1, instance2, ...
Statistic Name:	
MemoryUsageMB	The amount of memory in megabytes (MB) that is allocated to all Adaptive Server processes [Sequential = LST Non-Sequential = SUM] View Report: /report/sybaseASE/memory.rpt

Sybase ASE Server

Class:	Sybase ASE
Subclass:	Network
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Network
Open Table Name:	SYBASENET
Resource:	instance1, instance2, ...
Statistic Name:	
PacketErrors/s	The number of packet errors detected by the Adaptive Server per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/packets.rpt
PacketsReceived/s	The number of input packets read by the Adaptive Server per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/packets.rpt
PacketsSent/s	The number of output packets written by the Adaptive Server per second [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/packets.rpt
Class:	Sybase ASE
Subclass:	Procedure.Cache
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Procedure.Cache
Open Table Name:	SYBASEPROCEDURECACHE
Resource:	instance1, instance2, ...
Statistic Name:	
ProcedureReads/s	The number of procedure reads from disk during the interval [Sequential = AVG Non-Sequential = SUM]
ProcedureRecompiles/s	The number of times a procedure was executed with the recompile option in effect [Sequential = AVG Non-Sequential = SUM]
ProcedureRemovals/s	The number of times that a procedure aged out of cache [Sequential = AVG Non-Sequential = SUM]
ProcedureWrites/s	The number of procedure writes to disk during the interval [Sequential = AVG Non-Sequential = SUM]

Class:	Sybase ASE
Subclass:	SQL.Cache
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.SQL.Cache
Open Table Name:	SYBASESQLCACHE
Resource:	instance1, instance2, ...
Statistic Name:	
SqlStmtsCached	The number of SQL statements used in cache [Sequential = LST Non-Sequential = SUM]
SqlStmtsDroppedFromCache	The number of SQL statements removed from the cache [Sequential = LST Non-Sequential = SUM]
SqlStmtsInCache	The number of SQL statements in cache at the end of the interval [Sequential = LST Non-Sequential = SUM]
SqlStmtsNeverInCache	The number of SQL statements used but not eligible for cache [Sequential = LST Non-Sequential = SUM]
SqlStmtsNotInCache	The number of SQL statements no longer in cache at the end of the interval [Sequential = LST Non-Sequential = SUM]
SqlStmtsRestoredToCache	The number of SQL statements restored to cache [Sequential = LST Non-Sequential = SUM]
Class:	Sybase ASE
Subclass:	Sample
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Sample
Open Table Name:	SYBASESAMPLE
Resource:	instance1, instance2, ...
Statistic Name:	
tqsybase_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqsybase_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

Sybase ASE Server

Class:	Sybase ASE
Subclass:	System
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.System
Open Table Name:	SYBASESYS
Resource:	instance1, instance2, ...
Statistic Name:	
ActiveLocks	The number of active locks [Sequential = LST Non-Sequential = SUM] View Report: /report/sybaseASE/system.rpt
ConnectionsAttempted	The number of logins or attempted logins [Sequential = SUM Non-Sequential = SUM]
CurrentConnections	The number of connections at the time of the last sample [Sequential = SUM Non-Sequential = SUM] View Report: /report/sybaseASE/system.rpt
Databases	The number of databases that exist for the server [Sequential = LST Non-Sequential = SUM] View Report: /report/sybaseASE/system.rpt
Processes	The number of Adaptive Server processes [Sequential = LST Non-Sequential = SUM] View Report: /report/sybaseASE/system.rpt
Users	The number of Adaptive Server users [Sequential = LST Non-Sequential = SUM] View Report: /report/sybaseASE/system.rpt

Class:	Sybase ASE
Subclass:	Transaction Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Transaction Summary
Open Table Name:	SYBASETXNSUM
Resource:	instance1, instance2, ...
Statistic Name:	
APLClusteredTableInserts	The number of clustered table inserts with all of the pages locked [Sequential = SUM Non-Sequential = SUM]
APLHeapTableInserts	The number of heap table inserts with all of the pages locked [Sequential = SUM Non-Sequential = SUM]
TotalRowsDeleted	The number of table rows deleted [Sequential = SUM Non-Sequential = SUM]
TotalRowsUpdated	The number of table rows updated [Sequential = SUM Non-Sequential = SUM]
TranCount	The number of transactions committed [Sequential = SUM Non-Sequential = SUM]
TranRollbacks	The number of transaction rollbacks [Sequential = SUM Non-Sequential = SUM]

15.2. Active SQL Statistics

The Sybase ASE Agent (**tqsybase**) stores information about SQL processes that are currently active in the system. Statements with executions that do not cross over time intervals do not appear in this list.

Note: This table requires that the Sybase Monitor Data Analysis (MDA) tables be installed.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Active SQL
IT Resource Name:	/TeamQuest/System/systemname/Sybase ASE/instancename
TeamQuest Table Name:	Sybase ASE.Active SQL
Open Table Name:	SYBASEACTIVESQL
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
BatchID	The unique identifier for the SQL batch containing the SQL text [Sequential = ID Non-Sequential = ID]
ContextID	The stack frame of the procedure, if it is a procedure [Sequential = ID Non-Sequential = ID]
CpuTime	The amount of CPU time in seconds used by the statement [Sequential = SUM Non-Sequential = SUM]
Dbname	The database name [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Kpid	The kernel process identifier [Sequential = ID Non-Sequential = ID]
LogicalReads	The number of buffers read from the cache [Sequential = SUM Non-Sequential = SUM]
MemUsageMB	The amount of memory in megabytes (MB) that was used for execution of the statement [Sequential = SUM Non-Sequential = SUM]

NetworkPacketSize	The size in bytes of the network packet currently configured for the session [Sequential = ID Non-Sequential = ID]
NewPagesModified	The number of pages modified by the statement [Sequential = SUM Non-Sequential = SUM]
PacketsReceived	The number of network packets received by ASE [Sequential = ID Non-Sequential = ID]
PacketsSent	The number of network packets sent by ASE [Sequential = ID Non-Sequential = ID]
PhysicalReads	The number of buffers read from disk [Sequential = SUM Non-Sequential = SUM]
PlansAltered	The number of plans altered at execution time [Sequential = SUM Non-Sequential = SUM]
ProcedureID	The unique identifier for the procedure [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
ServerUserID	The server user identifier of the user executing the SQL [Sequential = ID Non-Sequential = ID]
Spid	The session process identifier [Sequential = ID Non-Sequential = ID]
SQLText	The SQL text [Sequential = ID Non-Sequential = ID]
StartTime	The date and time when the statement began execution [Sequential = SUM Non-Sequential = SUM]
System	The name of the system where the data is collected [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
WaitTime	The amount of time in seconds that the task has waited during execution of the statement [Sequential = SUM Non-Sequential = SUM]

15.3. Configuration Statistics

The Sybase ASE Agent (**tqsybase**) stores configuration data into a set of tables in the TeamQuest performance database. A record is stored only when a change in the configuration data is detected.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Configuration
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Configuration
Open Table Name:	SYBASECONF
Collection interval:	N/A
Default retention:	6 months
Table type:	Event

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Non-Sequential = ID]
ConfigNum	The configuration parameter number [Non-Sequential = ID]
ConfigValueChar	The value for the parameter datatype. The value is NULL for parameters with a datatype of integer. [Non-Sequential = ID]
ConfigValueInt	The value for the parameter with an integer datatype. The value is 0 for parameters with a datatype of character. This value can be modified by the user. [Non-Sequential = ID]
Default	The default value of the configuration parameter [Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained [Non-Sequential = ID]
Interval	The expected data sampling interval [Non-Sequential = ID]
MemoryUsedMB	The amount of memory in megabytes (MB) used by each configuration parameter [Non-Sequential = ID]
ParameterName	The name of the configuration parameter [Non-Sequential = ID]

RunValueChar	The current run value for the parameter with a datatype of character. The value is NULL for parameters with a datatype of integer. [Non-Sequential = ID]
RunValueInt	The current run value for the parameter with a datatype of integer. The value is 0 for the parameters with a datatype of character. [Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Non-Sequential = ID]
Shared Memory	Specifies if the memory used is shared memory. Values can be True, False, or <N/A> if the parameter does not use any memory. [Non-Sequential = ID]
System	The name of the system where the data is collected [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
Type	Specifies whether a configuration parameter is declared dynamic or static in its structure definition. The value can be one of the following: Dynamic (takes effect immediately) Static (takes effect after restarting Adaptive Server) [Non-Sequential = ID]
Unit	The unit of the parameter. The value can be one of the following: Not applicable—parameter has no units Number (number of items) Clock ticks (number of clock ticks) Microseconds Milliseconds Seconds Minutes Hours Days Bytes Kilobytes Megabytes Memory pages (2K) Virtual pages (2K) Logical pages Percent Ratio Switch (a Boolean value) ID (ID number) Name Rows [Non-Sequential = ID]

15.4. Database Detail Statistics

Detailed information about the Sybase ASE database statistics is stored within the Sybase ASE.Database Detail table in the TeamQuest performance database.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Database Detail
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Database Details
Open Table Name:	SYBASEDBDETAIL
Collection interval:	60 seconds (default)
Default retention:	10 days
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
CreateDate	The date the space was defined for the database [Sequential = ID Non-Sequential = ID]
Dbname	The database name [Sequential = ID Non-Sequential = ID]
Device_Fragments	The name of the device fragment [Sequential = ID Non-Sequential = ID]
FreeSizeMB	The available space in megabytes (MB) [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Owner	The user ID of the database owner [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
TotalSizeMB	The device size in megabytes (MB) [Sequential = ID Non-Sequential = ID]

15.5. Database Summary Statistics

Summary information about the Sybase ASE database statistics is stored within the Sybase ASE.Database Summary table in the TeamQuest performance database.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Database Summary
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Database Summary
Open Table Name:	SYBASEDBSUM
Collection interval:	60 seconds (default)
Default retention:	10 days
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
CreateDate	The date the space was defined for the database [Sequential = ID Non-Sequential = ID]
Dbname	The database name [Sequential = ID Non-Sequential = ID]
FreeSizeMB	The available space in megabytes (MB) [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Owner	The user ID of the database owner [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
TotalSizeMB	The database size in megabytes (MB) [Sequential = ID Non-Sequential = ID]

15.6. Device Detail Statistics

Detailed information about the Sybase ASE device statistics is stored within the Sybase ASE.Device Detail table in the TeamQuest performance database.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Device Detail
IT Resource Name:	/TeamQuest/System/systemname/Sybase ASE/instancename
TeamQuest Table Name:	Sybase ASE.Device Detail
Open Table Name:	SYBASEDEVDETAIL
Collection interval:	60 seconds (default)
Default retention:	10 days
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
DeviceName	The name of the storage device [Sequential = LST Non-Sequential = ID]
DeviceType	The identifier to determine if the device is defined as a Physical Disk, Dump Device, or Logical Disk [Sequential = LST Non-Sequential = ID]
FreeSizeMB	The available space in megabytes (MB) [Sequential = ID Non-Sequential = ID] View Report: /report/sybaseASE/diskutil.rpt
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
PctFreeSize	The percentage of disk space that is not being used [Sequential = ID Non-Sequential = ID]
PhysicalName	The system address of the storage device [Sequential = LST Non-Sequential = ID]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
TotalSizeMB	The device size in megabytes (MB) [Sequential = ID Non-Sequential = ID] View Report: /report/sybaseASE/diskutil.rpt

15.7. Engine Detail Statistics

Detailed information about the Sybase ASE engine statistics is stored within the Sybase ASE.Engine Detail table in the TeamQuest performance database.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Engine Detail
IT Resource Name:	/TeamQuest/System/systemname/Sybase ASE/instancename
TeamQuest Table Name:	Sybase ASE.Engine Detail
Open Table Name:	SYBASEENGINEDETAIL
Collection interval:	60 seconds (default)
Default retention:	10 days
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
BytesSent/Received	The total number of bytes received and sent for the engine [Sequential = SUM Non-Sequential = SUM]
ClockTicks	The number of clock ticks [Sequential = SUM Non-Sequential = SUM]
ContextSwitches	The number of times the Adaptive Server engine switched context from one user task to another [Sequential = SUM Non-Sequential = SUM]
CpuBusy	The amount of CPU time in seconds that the Adaptive Server CPU is doing Adaptive Server work for the engine [Sequential = SUM Non-Sequential = SUM]

EngineNum	The number of the engine [Sequential = ID Non-Sequential = ID]
EngineSleeps	The number of times the engine was idle [Sequential = SUM Non-Sequential = SUM]
Idle	The amount of CPU time in seconds that the Adaptive Server has been idle [Sequential = SUM Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
IoBusy	The amount of CPU time in seconds that the Adaptive Server has spent doing I/O operations [Sequential = SUM Non-Sequential = SUM]
PacketsSent/Received	The number of packets received and sent for the engine [Sequential = SUM Non-Sequential = SUM]
PctCpuBusy	The percentage of time the Adaptive Server is using the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/cputil.rpt
PctIdle	The percentage of time the Adaptive Server is not using the CPU [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/cputil.rpt
PctIoBusy	The percentage of time the Adaptive Server is doing I/O operations [Sequential = AVG Non-Sequential = SUM] View Report: /report/sybaseASE/cputil.rpt
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

15.8. Lock Detail Statistics

Detailed information about the Sybase ASE lock statistics is stored within the Sybase ASE.Lock Detail table in the TeamQuest performance database.

Note: This table requires that the Sybase Monitor Data Analysis (MDA) tables be installed.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Lock Detail
IT Resource Name:	/TeamQuest/System/systemname/Sybase ASE/instancename
TeamQuest Table Name:	Sybase ASE.Lock Detail
Open Table Name:	SYBASELOCKDETAIL
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Dbname	The database name [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Kpid	The kernel process identifier [Sequential = ID Non-Sequential = ID]
LockId	The lock object identifier [Sequential = ID Non-Sequential = ID]
LockLevel	The type of object for which the lock was requested. The value can be one of the following: Row Page Table Address [Sequential = ID Non-Sequential = ID]

LockState	Indicates whether the lock has been granted. The value can be one of the following: Hold Wait [Sequential = ID Non-Sequential = ID]
LockType	The type of lock. The value can be one of the following: Exclusive Shared Update [Sequential = ID Non-Sequential = ID]
PageNumber	The page that is locked when LockLevel = Page [Sequential = ID Non-Sequential = ID]
ParentSpid	The parent server process identifier [Sequential = ID Non-Sequential = ID]
RowNumber	The row that is locked when LockLevel = Row [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Spid	The server process identifier [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
WaitTime	The time in seconds that the lock request has not been granted [Sequential = SUM Non-Sequential = SUM]

15.9. Process Statistics

Process statistics are stored in the TeamQuest performance database tables by the Sybase ASE Agent (**tqsybase**). The process statistics are available when the processes have completed.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Process
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Sybase ASE/ <i>instancename</i>
TeamQuest Table Name:	Sybase ASE.Process
Open Table Name:	SYBASEPROC
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Affinity	The name of the engine to which the process has affinity [Sequential = ID Non-Sequential = ID]
Blocking_Spid	The server process identifier for blocking processes [Sequential = ID Non-Sequential = ID]
ClientHostName	The name by which the host is known for the current session [Sequential = ID Non-Sequential = ID]
ClientName	The name by which the user is known for the current session [Sequential = ID Non-Sequential = ID]
Cmd	The command currently being executed [Sequential = ID Non-Sequential = ID]
CPUSeconds	The CPU time the server application processes are using on the host. This number includes only Sybase Server internal processes, excluding the portion of the OS CPU time reported. [Sequential = SUM Non-Sequential = SUM]
Dbname	The database name [Sequential = ID Non-Sequential = ID]
EngineNum	The number of the engine on which the process is being executed [Sequential = ID Non-Sequential = ID]
Execlass	The execution class to which the process is bound [Sequential = ID Non-Sequential = ID]
GroupID	The group ID of the user who executed the command [Sequential = ID Non-Sequential = ID]

HostName	The name of the host computer [Sequential = ID Non-Sequential = ID]
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
IpAddress	The IP address of the client where the login is made [Sequential = ID Non-Sequential = ID]
Kpid	The kernel process identifier [Sequential = ID Non-Sequential = ID]
Lock_id	The lock owner identifier for a lock that is blocking a transaction [Sequential = ID Non-Sequential = ID]
LoggedInTime	The time and date when the client connected to Adaptive Server [Sequential = ID Non-Sequential = ID]
MemUsageMB	The amount of memory in megabytes (MB) allocated to the process [Sequential = LST Non-Sequential = ID]
NetworkPacketSize	The network packet size of the current connection [Sequential = ID Non-Sequential = ID]
Original_Server_Userid	The original server user identifier. If this field is not blank, the original server user identifier is used for the session authorization [Sequential = ID Non-Sequential = ID]
PhysicalIO	The number of accesses to the hard drive since each process started. This value includes accesses to the hard drive for physical reads and physical writes. [Sequential = SUM Non-Sequential = SUM]
Priority	The base priority associated with the process [Sequential = ID Non-Sequential = ID]
ProgramName	The name of the front-end module [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Server_Userid	The server user identifier of the user executing the process [Sequential = ID Non-Sequential = ID]
Spid	The server process identifier [Sequential = ID Non-Sequential = ID]

Status	<p>The Process ID status. The value can be one of the following:</p> <ul style="list-style-type: none"> Infected Background Recv sleep Send sleep Alarm sleep Lock sleep Sleeping Runnable Running Stopped Bad status Log suspend <p>[Sequential = LST Non-Sequential = ID]</p>
System	<p>The name of the system where the data is collected</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
Time_Blocked	<p>The amount of time being blocked in seconds</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Tranname	<p>The name of the active transaction</p> <p>[Sequential = LST Non-Sequential = ID]</p>
Userid	<p>The ID of the user who executed the command</p> <p>[Sequential = ID Non-Sequential = ID]</p>

15.10. Process Waits Statistics

The Sybase ASE.Process Waits table stores detailed information about current wait events by a process.

Note: This table requires that the Sybase Monitor Data Analysis (MDA) tables be installed.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	Process Waits
IT Resource Name:	/TeamQuest/System/systemname/Sybase ASE/instancename
TeamQuest Table Name:	Sybase ASE.Process Waits
Open Table Name:	SYBASEPROCWAITS
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
AvgWaitTime	The average amount of time in seconds the process spent waiting for the event [Sequential = AVG Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
Kpid	The kernel process identifier [Sequential = ID Non-Sequential = ID]
LoginId	The login name for the process that is waiting [Sequential = ID Non-Sequential = ID]
ProcessWaits	The number of times the process has waited for the event [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Spid	The server process identifier [Sequential = ID Non-Sequential = ID]
System	The name of the system where the data is collected [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
WaitEventID	The unique identifier for the wait event [Sequential = ID Non-Sequential = ID]

15.11. System Waits Statistics

The Sybase ASE.System Waits table stores detailed information about current wait events by a system.

Note: This table requires that the Sybase Monitor Data Analysis (MDA) tables be installed.

Table Field Hierarchy

Class:	Sybase ASE
Subclass:	System Waits
IT Resource Name:	/TeamQuest/System/systemname/Sybase ASE/instancename
TeamQuest Table Name:	Sybase ASE.System Waits
Open Table Name:	SYBASESYSWAITS
Collection interval:	60 seconds (default)
Default retention:	1 day
Table type:	Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
AvgWaitTime	The average amount of time in seconds the process spent waiting for the event [Sequential = SUM Non-Sequential = SUM]
Instance	The name of the instance from which the data is obtained [Sequential = ID Non-Sequential = ID]
Interval	The expected data sampling interval [Sequential = SUM Non-Sequential = ID]
LoginId	The ID of the user who executed the command [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
WaitEventID	The unique identifier for the wait event [Sequential = ID Non-Sequential = ID]
Waits	The number of times tasks have waited for the event [Sequential = SUM Non-Sequential = SUM]

Section 16

System Alarm Statistics

Alarms generated by TeamQuest Manager are recorded in the System Alarms table. This table can be managed and viewed using the TeamQuest performance applications (TeamQuest Analyzer, TeamQuest View, and TeamQuest Alert).

The System Alarms table is created by the Alarm Service (**tqalm**). The table is located in your database directory *TQDATADIR\dbname\alarms*.

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

System Alarm Statistics

Table Field Hierarchy

Class:	System
Subclass:	Alarms
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	System.Alarms
Open Table Name:	SYSALARMS
Collection interval:	N/A
Default retention:	1 month
Table type:	Event

Statistic Name	Description
----------------	-------------

The combination of Sequence_Number, System, and Time fields is the unique identifier for an alarm within TeamQuest Manager.

Alarm_Status	The timestamp of the data sample. This is a hidden field and is for internal use only. [Non-Sequential = ID]
Sequence_Number	A sequential number assigned to the alarm for uniqueness [Non-Sequential = ID]
System	The name of the system for which the alarm occurred. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The date and time at which the alarm occurred. When the Alarm_Type is 1, the Event_Timestamp field may be a more accurate indicator of when the alarm actually occurred. [Non-Sequential = ID]

The combination of `Original_Timestamp`, `Original_Sequence_Number`, and `Original_System` is the unique identifier for when the alarm was last raised. For the first time the alarm is raised, or after a previous alarm clear, these have the same values as `Sequence_Number`, `System`, and `Time`. As the severity escalates on an alarm, these fields will reflect when the alarm was last raised. For example, if an alarm goes from Warning to Critical, these fields on the Critical alarm record will contain the values of `Sequence_Number`, `System`, and `Time` from the Warning alarm record. These fields can also be used to match up an alarm clear with when the alarm was raised.

Alarm_Conditions A string that shows the conditions that caused the alarm to be generated (and values of any aliases referenced only in the alarm action). The last condition is followed by a trailing null instead of a new line. For example:

```
alias1(45.40) > 25.00 and < 50
alias2(95.50) > 85.00
```

The value inside the parentheses is the current value for the parameter, derived statistic, or table field referenced by the alias. The mathematical operators and threshold values will only be present when an alarm is being raised.
[Non-Sequential = ID]

Alarm_ID The user-defined alarm identifier assigned to the alarm. It will also contain specific resource and workload names if needed.
[Non-Sequential = ID]

Alarm_Message A user-defined message from the alarm action definition that contains information regarding the alarm. If no information exists, the field will be blank. This message can be sent as a console message by updating Console Message setting to ON in the Alarm Service configuration file of TeamQuest Manager.
[Non-Sequential = ID]

Alarm_Type This has a value of 1 for an event alarm and 0 for alarms that have both a raised and a clear condition. This is a hidden field and is for internal use only.
[Non-Sequential = ID]

Event_Timestamp The actual time that the event occurred. This is always the same as the `Time` field in the alarm record unless table data is used in the alarm. When using table data, this field will be set to the value of the `Time` field in the table record where the alarm was detected. If there is no `Time` field in the table record, the `Event_Timestamp` field will be the same as the `Time` field of the alarm record.
[Non-Sequential = ID]

Count A count of alarms that is useful when consolidating the alarms at report time
[Non-Sequential = SUM]

Database_Name The name of the performance database and aggregation set used to detect the alarm condition. The names are separated by a colon (Database:Aggset).
[Non-Sequential = ID]

System Alarm Statistics

Parameter_Information	<p>A string made up of all alias names and the associated parameters, derived statistics, or table field names that are referenced in the alarm definition.</p> <p>[Non-Sequential = ID]</p> <p>The parameters, derived statistics, and table field names are colon-delimited and the last one is followed by a trailing null instead of a new line. Table and parameter aliases are grouped together and have a section name to identify them. If there are no section names present, all aliases are assumed to be parameter aliases.</p> <p>For example:</p> <pre>[Parameter] conn1 = saturn:Web Server:Summary::connections/sec:WebServer1:: conn2 = saturn:Web Server:Summary::connections/sec:WebServer2:: [Table] login = NT:Process:login cmd = NT:Process:command</pre>
Original_Timestamp	<p>The timestamp when the alarm was last raised</p> <p>[Non-Sequential = ID]</p>
Original_Sequence_Number	<p>The Sequence_Number when the alarm was last raised</p> <p>[Non-Sequential = ID]</p>
Original_System	<p>The system where the alarm was last raised</p> <p>[Non-Sequential = ID]</p>
Sample_Interval	<p>The frequency of how often the Alarm Service checks for alarms (in seconds). The value in this field is the same as the sample rate of the aggregation set specified in the Database_Name field. This is a hidden field and is for internal use only.</p> <p>[Non-Sequential = LST]</p>
Severity	<p>The severity of the alarm (Normal, Warning, Minor, Major, or Critical)</p> <p>[Non-Sequential = ID]</p>
Severity_Number	<p>The alarm severity as a numeric value:</p> <p>0 = Unknown 1 = Normal 2 = Warning 3 = Minor 4 = Major 5 = Critical</p> <p>[Non-Sequential = ID]</p>
SNMP_Message	<p>The text of the SNMP message generated for the alarm. If no text exists, the field will be blank.</p> <p>[Non-Sequential = ID]</p>

Table_Key	A string containing the primary key information pointing to the table record that caused the alarm to be raised or cleared for alarms that use table data. It is a set of Fieldname = Value pairs. The last pair is followed by a trailing null instead of a new line. If table data is not used in the alarm, this is an empty string. [Non-Sequential = ID]
Table_Name	A string containing the class and name of the table that is used in the alarm (Class:Name) for alarms that use table data. If table data is not used in the alarm, this is an empty string. [Non-Sequential = ID]

Section 17

VMware Systems

The VMware Infrastructure Agent (**tqvmwarep**) remotely collects performance data for VMware vCenter servers, VMware hosts, and VMware virtual machines.

This section contains a listing of the statistics collected by the agent:

- Physical Tables (see 17.1)
- Derived Tables (see 17.2)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

17.1. Physical Tables

A physical table definition allows multiple stored tables to be brought together into a single logical table. A derived table is created by indicating which fields from one or more stored tables are brought together, and which fields are to be calculated based on values of other fields in the same record when the stored tables have been brought together.

A derived table definition can reference one or more stored tables. A derived table cannot reference another derived table. One of the stored tables referenced by a derived table must be identified as the *primary reference table*. The primary reference table is the table that the other tables are joined to. The other reference tables are referred to as the *secondary reference tables*. One of the reference tables can be flagged as the table to use when determining the selection statements for applying an IT Resource to the derived table. This table is referred to as the *primary selection table*.

In this subsection, you can find a listing of the physical table statistics collected by the agent:

- Block Device.by Host System Device Table (see 17.1.1)
- Block Device.by Virtual Machine Table (see 17.1.2)
- Block Device.VMware Summary Table (see 17.1.3)
- CPU.by Host Processor Table (see 17.1.4)
- CPU.by Virtual CPU Table (see 17.1.5)
- CPU.by Virtual Machine Table (see 17.1.6)
- CPU.by VMware Resource Table (see 17.1.7)
- CPU.Relative Performance Table (see 17.1.8)
- CPU.VMware Summary Table (see 17.1.9)
- HINV.CPUModel Table (see 17.1.10)
- HINV.CPU Thread Speeds Table (see 17.1.11)
- HINV.Devices Table (see 17.1.12)
- HINV.FileSystem Table (see 17.1.13)
- HINV.Summary Table (see 17.1.14)
- Memory.by Virtual Machine Table (see 17.1.15)
- Memory.VMware Summary Table (see 17.1.16)
- Network Device.by Host System Device Table (see 17.1.17)
- Network Device.by Virtual Machine Table (see 17.1.18)
- Network Device.vmnic by Virtual Machine Table (see 17.1.19)
- Network Device.VMware Summary Table (see 17.1.20)
- VMware.Availability by Virtual Machine Table (see 17.1.21)

- VMware.Host Configuration Table (see 17.1.22)
- VMware.Host Status Table (see 17.1.23)
- VMware.Storage Configuration Table (see 17.1.24)
- VMware.Support Metrics Table (see 17.1.25)
- VMware.Virtual_Machines Table (see 17.1.26)
- VMware Cluster.CPU Summary Table (see 17.1.27)
- VMware Cluster.Memory Summary Table (see 17.1.28)
- VMware Cluster.Resource Allocation Table (see 17.1.29)
- VMware Cluster.Virtual Machine Operations (see 17.1.30)
- VMware Datastore.File Type Usage by Datacenter Table (see 17.1.31)
- VMware Datastore.Summary Table (see 17.1.32)
- VMware Datastore.Usage by Virtual Machine Table (see 17.1.33)
- VMware Resource Pool.CPU Summary Table (see 17.1.34)
- VMware Resource Pool.Memory Summary Table (see 17.1.35)
- VMware Resource Pool.Resource Allocation Table (see 17.1.36)
- VMware Storage.Adapter by Host System Table (see 17.1.37)
- VMware Storage.Adapter Summary Table (see 17.1.38)
- VMware Storage.Datastore by Host System (see 17.1.39)
- VMware Storage.Datastore by Virtual Machine (see 17.1.40)
- VMware Storage.Datastore Summary (see 17.1.41)
- VMware Storage.Path by Host System Table (see 17.1.42)
- VMware Storage.Path Summary Table (see 17.1.43)
- VMware Storage.Virtual Disk by Virtual Machine Table (see 17.1.44)

17.1.1. Block Device.by Host System Device Table

The Block Device.by Host System Device table stores device-level I/O data for all of the logical unit numbers associated with VMware hosts.

Table Field Hierarchy

Class:	Block Device
Subclass:	by Host System Device
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	Block Device.by Host System Device
Open Table Name:	BLKDEVBYHOSTSYSDEVIC
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Host Block Device Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgDeviceLatency	The average amount of time in milliseconds taken to complete a command to the physical device [Sequential = AVG Non-Sequential = AVG]
avgDeviceReadLatency	The average amount of time in milliseconds taken to complete a read from the physical device [Sequential = AVG Non-Sequential = AVG]
avgDeviceWriteLatency	The average amount of time in milliseconds taken to complete a write to the physical device [Sequential = AVG Non-Sequential = AVG]
avgKernelLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per command [Sequential = AVG Non-Sequential = AVG]
avgKernelReadLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per read [Sequential = AVG Non-Sequential = AVG]
avgKernelWriteLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per write [Sequential = AVG Non-Sequential = AVG]

avgLatency	The average amount of time in milliseconds taken to complete a command request (queue and disk service time) by the host system disk [Sequential = AVG Non-Sequential = AVG]
avgQueueLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per command [Sequential = AVG Non-Sequential = AVG]
avgQueueReadLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per read [Sequential = AVG Non-Sequential = AVG]
avgQueueWriteLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per write [Sequential = AVG Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of a guest operating system [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of a guest operating system [Sequential = AVG Non-Sequential = AVG]
busRst/s	The number of bus resets per second that occurred on the host system disk [Sequential = AVG Non-Sequential = SUM]
cmds/s	The number of commands (requests) issued per second to the host system disk [Sequential = AVG Non-Sequential = SUM]
cmdsAbrt/s	The number of commands aborted by the host system disk per second [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the host system disk [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the host system disk [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data transferred per second in kilobytes (KB) by the host system disk [Sequential = AVG Non-Sequential = SUM]
maxQueueDepth	The maximum queue depth. The maximum number of I/O operations supported by the LUN that can be outstanding at a given time. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]

reads/s	The number of read requests issued per second to the host system disk [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the disk device [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totalTime	The total time in milliseconds for all command requests on a VMware host [Sequential = SUM Non-Sequential = SUM]
writes/s	The number of write requests issued per second to the host system disk [Sequential = AVG Non-Sequential = SUM]

17.1.2. Block Device.by Virtual Machine Table

The Block Device.by Virtual Machine table stores device-level I/O data for disk devices from the perspective of the virtual machine.

The Block Device.by Virtual Machine table data includes VMware host management agent data, which contains the consumption of disk resources by the VMware host. The virtual machine name of this record is the same as the VMware host name.

Table Field Hierarchy

Class:	Block Device
Subclass:	by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Block Device.by Virtual Machine
Open Table Name:	BLKDEVBYVM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Virtual Machine Block Device Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgDeviceLatency	The average amount of time in milliseconds taken to complete a command to the physical device [Sequential = AVG Non-Sequential = AVG]
avgDeviceReadLatency	The average amount of time in milliseconds taken to complete a read from the physical device [Sequential = AVG Non-Sequential = AVG]
avgDeviceWriteLatency	The average amount of time in milliseconds taken to complete a write to the physical device [Sequential = AVG Non-Sequential = AVG]
avgKernelLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per command [Sequential = AVG Non-Sequential = AVG]
avgKernelReadLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per read [Sequential = AVG Non-Sequential = AVG]

avgKernelWriteLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per write [Sequential = AVG Non-Sequential = AVG]
avgLatency	The average amount of time in milliseconds taken to complete a command request (queue and disk service time) by the host system disk [Sequential = AVG Non-Sequential = AVG]
avgQueueLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per command [Sequential = AVG Non-Sequential = AVG]
avgQueueReadLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per read [Sequential = AVG Non-Sequential = AVG]
avgQueueWriteLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per write [Sequential = AVG Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of a guest operating system [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of a guest operating system [Sequential = AVG Non-Sequential = AVG]
busRst/s	The number of bus resets per second that occurred on the virtual machine [Sequential = AVG Non-Sequential = SUM]
cmds/s	The number of commands (requests) issued per second to the physical device [Sequential = AVG Non-Sequential = SUM]
cmdsAbrt/s	The number of commands per second that were aborted by the virtual machine [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data transferred per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]

maxTotalLatency	The highest latency value across all disks used by the host. Latency measures the time taken to process a SCSI command issued by the guest OS to the virtual machine. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = MAX Non-Sequential = MAX]
reads/s	The number of read requests issued per second to the physical device [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the physical device [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totalTime	The total time in milliseconds for all command requests by the virtual machine [Sequential = SUM Non-Sequential = SUM]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
writes/s	The total number of write requests issued per second to the physical device [Sequential = AVG Non-Sequential = SUM]

17.1.3. Block Device.VMware Summary Table

The Block Device.VMware Summary table stores device-level I/O data for disk devices from the perspective of the VMware host.

Table Field Hierarchy

Class:	Block Device
Subclass:	VMware Summary
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	Block Device.VMware Summary
Open Table Name:	BLKDEVVMSUM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Host Block Device Summary

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
cmds/s	The total number of read and write command requests per second [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data read and written per second in kilobytes (KB) for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]
maxTotalLatency	The highest latency value across all disks used by the host. Latency measures the time taken to process a SCSI command issued by the guest OS to the virtual machine. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = MAX Non-Sequential = MAX]

reads/s	The total number of read requests per second for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
writes/s	The total number of write requests per second for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]

17.1.4. CPU.by Host Processor Table

The CPU.by Host Processor table stores device-level utilization data for processors from the perspective of the VMware host system.

Table Field Hierarchy

Class:	CPU
Subclass:	by Host Processor
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	CPU.by Host Processor
Open Table Name:	CPUBYHOSTPROC
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Host CPU Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
idle	The amount of processor time in seconds that is spent in an idle state [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Object	The name of the CPU object [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

usage	The percentage of time the CPU is in use over the collection interval [Sequential = AVG Non-Sequential = SUM]
usedsec	The processor time consumed by the VMware host in seconds [Sequential = SUM Non-Sequential = SUM]

17.1.5. CPU.by Virtual CPU Table

The CPU.by Virtual CPU table stores utilization data from virtual processors from the perspective of the virtual machine.

Table Field Hierarchy

Class:	CPU
Subclass:	by Virtual CPU
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	CPU.by Virtual CPU
Open Table Name:	CPUBYVIRTCPU
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Virtual Machine Virtual CPU Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
maxlimited	The amount of time in seconds the virtual machine was ready to run but did not run because it reached the maximum CPU limit setting. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
Object	The name of the CPU object [Sequential = ID Non-Sequential = ID]
overlap	The amount of time in seconds the virtual machine was interrupted while performing system services. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]

ready	The virtual CPU time that is spent in the ready state in seconds [Sequential = SUM Non-Sequential = SUM]
run	The total amount of time in seconds scheduled for the CPU by the virtual CPU. This time does not account for hyper-threading and system time. On a hyper-threading enabled server, the percentage of run time can be twice as large as the percentage of CPU used time. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
syssec	The virtual CPU time that is spent on system processes in seconds [Sequential = SUM Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
usage(MHz)	The CPU usage in megahertz (MHz) over the collection interval [Sequential = AVG Non-Sequential = SUM]
usedsec	The virtual CPU time that is used in seconds [Sequential = SUM Non-Sequential = SUM]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
waitsec	The virtual CPU wait time in seconds [Sequential = SUM Non-Sequential = SUM]

17.1.6. CPU.by Virtual Machine Table

The CPU.by Virtual Machine table stores utilization data for processors summarized to the virtual machine level.

The CPU.by Virtual Machine table data includes VMware host management agent data, which contains the consumption of CPU resources by the VMware host. The virtual machine name of this record is the same as the VMware host name.

Table Field Hierarchy

Class:	CPU
Subclass:	by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	CPU.by Virtual Machine
Open Table Name:	CPUBYVM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Virtual Machine CPU Usage

Statistic Name	Description
%busy	The percentage of the server processor or processors that the virtual machine used [Sequential = AVG Non-Sequential = SUM]
%vcpu_busy	The percentage of the virtual machine's virtual processors used [Sequential = AVG Non-Sequential = SUM]
%vcpu_ready	The percentage of time the virtual machine was ready to perform an operation but had to wait for a processor [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
costop	The amount of time in seconds a Symmetric Multi-Processing (SMP) virtual machine was ready to run, but was delayed due to console virtual CPU (co-vCPU) scheduling contention. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]

demand	The rate of CPU demand in megahertz (MHz). This value represents the average active CPU load in the sampling interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
entitlement	The share of the CPU resource in megahertz (MHz) that a virtual machine should get as a result of the virtual CPU count and assigned shares. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
host_uptime	The elapsed time in seconds between two samples that the host or virtual machine was powered on [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
latency	The percentage of time the virtual machine was ready to run but was not scheduled to run because of physical CPU resource contention. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
ready	The amount of time in seconds the virtual machine was ready to perform an operation but had to wait for a processor [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shares	The number of CPU shares allocated to the virtual machine [Sequential = LST Non-Sequential = SUM]
swapwait	The amount of time in seconds the virtual machine was waiting for swap page-ins. CPU swap wait is included in CPU wait. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
syssec	The amount of system time in seconds consumed by the virtual machine [Sequential = SUM Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totalCapacity	Total amount of CPU reservation (in megaHertz) used by and available for powered-on virtual machines and vSphere services on the host. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
uptime	The elapsed time in seconds between two samples that the host or virtual machine was powered on [Sequential = SUM Non-Sequential = SUM]
usage(MHz)	The CPU usage in megahertz (MHz) over the collected interval [Sequential = AVG Non-Sequential = SUM]

usedsec	The processor time in seconds consumed by the virtual machine [Sequential = SUM Non-Sequential = SUM]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
waitsec	The virtual CPU wait time in seconds [Sequential = SUM Non-Sequential = SUM]

17.1.7. CPU.by VMware Resource Table

The CPU.by VMware Resource table stores processor utilization data on how the hypervisor is managing the CPU.

Table Field Hierarchy

Class:	CPU
Subclass:	by VMware Resource
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	CPU.by VMware Resource
Open Table Name:	CPUBYVMRES
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Host CPU Resource Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Object	The name of the object for the host system [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
usage(MHz)	The CPU usage in megahertz (MHz) over the collection interval [Sequential = AVG Non-Sequential = SUM]

17.1.8. CPU.Relative Performance Table

The CPU.Relative Performance table stores the current relative performance data for physical systems.

Table Field Hierarchy

Class:	CPU
Subclass:	Relative Performance
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	CPU.Relative Performance
Open Table Name:	CPURELPERF
Collection interval:	1 minute
Default retentions:	1 month
Table type:	Performance
Derived tables using fields from this table:	VMware CPU Relative Performance

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
relative_unused	The amount of CPU resources not used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
rel_used	The amount of CPU resources used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]

System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

17.1.9. CPU.VMware Summary Table

The CPU.VMware Summary table stores CPU performance data summarized by VMware hosts.

Table Field Hierarchy

Class:	CPU
Subclass:	VMware Summary
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	CPU.VMware Summary
Open Table Name:	CPUVMSUM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Host CPU Summary

Statistic Name	Description
%busy	The percentage of the CPU used [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
costop	The amount of time in seconds a Symmetric Multi-Processing (SMP) virtual machine was ready to run, but was delayed due to console virtual CPU (co-vCPU) scheduling contention. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
demand	The rate of CPU demand in megahertz (MHz). This value represents the average active CPU load in the sampling interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]

idle	The processor time in seconds that is spent in an idle state [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
latency	The percentage of time the host was ready to run but was not scheduled to run because of CPU resource contention. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
online_cpus	The number of logical CPUs that were online [Sequential = LST Non-Sequential = SUM]
online_cpus_physical	The number of physical CPUs that were online [Sequential = LST Non-Sequential = SUM]
reservedCapacity	The total CPU capacity in megahertz (MHz) reserved by all of the virtual machines [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
swapwait	The amount of time in seconds the virtual machine was waiting for swap page-ins. CPU swap wait is included in CPU wait. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totalCapacity	The total amount of CPU reservation in megahertz (MHz) used by and available for powered-on virtual machines and VMware vSphere services on the host. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
uptime_t	The total time in days elapsed since the last VMware host reboot [Sequential = LST Non-Sequential = SUM]
usage(MHz)	The CPU usage in megahertz (MHz) over the collected interval [Sequential = AVG Non-Sequential = SUM]

17.1.10. HINV.CPUModel Table

The HINV.CPUModel table stores best-match, relative performance data about the system configuration. This table is created by the VMware Infrastructure Agent (**tgvmwarep**) to map physical hardware to a CPU model that describes performance in relative terms. This table is not created for any virtualized system. It is populated for physical systems and VMware hosts. It is not populated for VMware guests, Hyper-V guests, Solaris logical domains (LDOMs), Solaris guest LDOMs, KVM guests, or Linux on POWER systems.

Table Field Hierarchy

Class:	HINV
Subclass:	CPUModel
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.CPUModel
Open Table Name:	HINVCPU
Collection interval:	N/A
Default retention:	1 year
Table type:	State
Derived tables using fields from this table:	N/A

Statistic Name	Description
cpu_chips	The number of CPU chips or sockets [Non-Sequential = ID]
cpu_confidence	The percentage of confidence in the correctness of the CPU match based on model, frequency, and configuration (chips, cores, threads) [Non-Sequential = ID]
cpu_cores	The number of CPU cores or processors on an individual CPU chip [Non-Sequential = ID]
cpu_name	The name of the selected CPU [Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Non-Sequential = ID]
cpu_speed	The speed of the processor in megahertz (MHz) or gigahertz (GHz) [Non-Sequential = ID]
cpu_threads	The number of CPU threads on an individual CPU core or processor [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]

Time	The timestamp of the data sample [Non-Sequential = ID]
user_override	The user override status of the default TeamQuest generated CPU match. This field is not currently used and should appear as 0. [Non-Sequential = ID]

17.1.11. HINV.CPU Thread Speeds Table

The HINV.CPU Thread Speeds table stores best-match, performance improvement factors based on the number of active threads per core.

Table Field Hierarchy

Class:	HINV
Subclass:	CPU Thread Speeds
IT Resource Name:	/TeamQuest/System/systemname
TeamQuest Table Name:	HINV.CPU Thread Speeds
Open Table Name:	HINVCPUTHREADSPEEDS
Collection interval:	N/A
Default retention:	1 year
Table type:	State

Statistic Name	Description
speed_up_factor	The performance improvement when there are multiple active threads per core, compared to when there is only one active thread per core [Sequential = AVG Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
thread_number	The number of active threads [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = ID Non-Sequential = ID]

17.1.12. HINV.Devices Table

The HINV.Devices table stores configuration data for physical block devices and respective device controllers in VMware hosts.

Table Field Hierarchy

Class:	HINV
Subclass:	Devices
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	HINV.Devices
Open Table Name:	HINVDEVS
Collection interval:	N/A
Default retention:	1 year
Table type:	State
Derived tables using fields from this table:	N/A

Statistic Name	Description
class	The device classification. For example, controller, disk, tape, or cdrom. [Non-Sequential = ID]
controller	The device path indicator which defines a connection to another device [Non-Sequential = ID]
lun_id	The globally unique Logical Unit Number (LUN) identifier for Storage Area Network (SAN) based disk devices. The LUN identifier can be a Network Address Authority (NAA), Extended Unique Identifier (EUI), or iSCSI Qualified Name (IQN) value. This field is reported as <N/A> for non-SAN based disk devices, CD-ROM drives, tape drives, and so on. [Non-Sequential = ID]
name	The unique identifier for this device [Non-Sequential = ID]
name2	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = ID]
product	The product identifier. This field may be blank. [Non-Sequential = ID]
revision	The revision level for this product. This field may be blank. [Non-Sequential = ID]
rpm	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = ID]
sequence	The sequence number of the device [Non-Sequential = ID]

swap	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
vendor	The name of the device vendor. This field may be blank. [Non-Sequential = ID]

17.1.13. HINV.FileSystem Table

The HINV.FileSystem table stores configuration data about host file systems.

Table Field Hierarchy

Class:	HINV
Subclass:	FileSystem
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	HINV.FileSystem
Open Table Name:	HINVFILESYS
Collection interval:	N/A
Default retention:	1 year
Table type:	State
Derived tables using fields from this table:	N/A

Statistic Name	Description
BlkSize	The size of a block on the file system in bytes [Non-Sequential = ID]
Device	The path for the device on which the file system is mounted [Non-Sequential = ID]
Name	The unique identifier for the file system [Non-Sequential = ID]
Source	The source physical disk or logical volume of the file system. This field is always blank for this platform. [Non-Sequential = ID]
System	The name by which the system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]

TotBlks	The total number of blocks on the file system [Non-Sequential = ID]
TotFiles	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = ID]
TotSize	The total amount of space on the file system in megabytes [Non-Sequential = ID]
Type	The type of the file system [Non-Sequential = ID]

17.1.14. HINV.Summary Table

The HINV.Summary table stores configuration data about VMware hosts.

Table Field Hierarchy

Class:	HINV
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	HINV.Summary
Open Table Name:	HINVSUM
Collection interval:	N/A
Default retention:	1 year
Table type:	State
Derived tables using fields from this table:	N/A

Statistic Name	Description
core_multi-thread	The status or ability of the processor to support multiple independent threads. This field contains <N/A> if the information is not available. For VMware guests, this field contains <N/A>. [Non-Sequential = ID]
cores_per_chip	The number of cores or processors on an individual chip. This field contains <N/A> if the information is not available. For VMware guests, this field contains <N/A>. [Non-Sequential = ID]
cpu_chips	The number of CPU chips or sockets. This field contains <N/A> if the information is not available. For VMware guests, this field contains <N/A>. [Non-Sequential = ID]
cpu_count	The number of configured processors [Non-Sequential = ID]
cpu_speed	The speed of the processor in MHz or GHz [Non-Sequential = ID]

cpu_type	The basic instruction set architecture of the current system [Non-Sequential = ID]
logical_cpu_count	The number of logical processors [Non-Sequential = ID]
mem_size	The size of configured random access memory in kilobytes, where 1 kilobyte = 1,024 bytes [Non-Sequential = ID]
memory	The size of configured random access memory in megabytes [Non-Sequential = ID]
memory_size	The size of configured random access memory in megabytes or gigabytes [Non-Sequential = ID]
model	The name of the hardware implementation or platform [Non-Sequential = ID]
os_release	The name and level of the implementation of the operating system [Non-Sequential = ID]
pagesize	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = ID]
partition_type	The partition type of the system. This field will contain VMware Host if the system is a VMware host system. [Non-Sequential = ID]
serial	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = ID]
System	The name by which the host system is known to a communication network or node name. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
system_identifier	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = ID]
system_type	The name of the operating system [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = LST]
timezone	The time zone where the data was collected [Non-Sequential = ID]
TQLevel	The level of TeamQuest Manager [Non-Sequential = ID]

17.1.15. Memory.by Virtual Machine Table

The Memory.by Virtual Machine table stores performance data related to memory for VMware virtual machines.

The Memory.by Virtual Machine table data includes VMware host management agent data, which contains the consumption of memory resources by the VMware host. The virtual machine name of this record is the same as the VMware host name.

Table Field Hierarchy

Class:	Memory
Subclass:	by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Memory.by Virtual Machine
Open Table Name:	MEMBYVM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Virtual Machine Memory Usage

Statistic Name	Description
%usage	The percentage of total available memory that is used [Sequential = AVG Non-Sequential = SUM]
active	The working set size estimate in megabytes at the end of the interval for the virtual machine [Sequential = AVG Non-Sequential = SUM]
activewrite	The amount of memory in megabytes actively being written to or by the virtual machine. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
consumed	The amount of host memory in megabytes consumed by the virtual machine for guest memory [Sequential = AVG Non-Sequential = SUM]
entitlement	The share of the memory resource in megabytes that a virtual machine should get as a result of the virtual CPU count and assigned shares. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]

Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
latency	The percentage of time the virtual machine was waiting to access swapped or compressed memory. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
lowfreethreshold	The threshold of free host physical memory in megabytes. The VMware ESX Server will begin reclaiming memory from virtual machines through ballooning and swapping if the amount of free host physical memory falls below this value. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
memctl	The amount of memory in megabytes currently reclaimed using vmmemctl for the virtual machine [Sequential = AVG Non-Sequential = SUM]
memctlgt	The target memory size in megabytes to reclaim using vmmemctl for the virtual machine [Sequential = AVG Non-Sequential = SUM]
overheadMax	The amount of memory in megabytes reserved for the virtualization overhead for the virtual machine. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
overheadTouched	The amount of actively touched overhead memory in megabytes reserved for the virtualization overhead for the virtual machine. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shared	The amount of memory in megabytes shared between all running virtual machines and within a virtual machine [Sequential = AVG Non-Sequential = SUM]
shares	The number of memory shares allocated to the virtual machine [Sequential = LST Non-Sequential = SUM]
size	The amount of memory in megabytes currently allocated to the virtual machine [Sequential = AVG Non-Sequential = SUM]
swapin	The total amount of memory in megabytes that has been read from the virtual machine's swap file to the machine memory by the VMKernel during the interval [Sequential = AVG Non-Sequential = SUM]

swapinRate	<p>The rate memory is swapped from the disk into active memory during the collection interval. This value applies to virtual machines. This value is more useful than the swapin statistic to determine if the virtual machine is running slow due to memory swapping, especially when evaluating real-time data. This statistic is available for VMware ESX 4.0.0 and later.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
swapout	<p>The total amount of memory in megabytes that has been transferred the virtual machine's swap file to the machine memory by the VMKernel during the interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
swapoutRate	<p>The rate memory is swapped from active memory to the disk during the collection interval. This value applies to virtual machines. This value is more useful than the swapout statistic to determine if the virtual machine is running slow due to memory swapping, especially when evaluating real-time data. This statistic is available for VMware ESX 4.0.0 and later.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
swapped	<p>The amount of memory in megabytes currently swapped to the VMware File System 3 (VMFS3) swap file</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
swaptgt	<p>The target size in megabytes to swap to the VMware File System 3 (VMFS3) swap file for the virtual machine</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
System	<p>The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
totalCapacity	<p>The total amount of memory reservation in megabytes used by and available for powered-on virtual machines and VMware vSphere services on the host. This statistic is available for VMware ESX 4.1.0 and later.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Virtual_Machine	<p>The name of the virtual machine to which the data applies</p> <p>[Sequential = ID Non-Sequential = ID]</p>
zero	<p>The amount of memory in megabytes that is zeroed out</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>

zipped	The amount of zipped memory in megabytes. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = LST Non-Sequential = SUM]
zipSaved	The amount of memory in megabytes saved due to memory zipping. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = LST Non-Sequential = SUM]

17.1.16. Memory.VMware Summary Table

The Memory.VMware Summary table stores memory data for VMware host systems.

Table Field Hierarchy

Class:	Memory
Subclass:	VMware Summary
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	Memory.VMware Summary
Open Table Name:	MEMVMSUM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Host Memory Summary

Statistic Name	Description
%usage	The percentage of memory usage over the collection interval [Sequential = AVG Non-Sequential = SUM]
active	The working set size estimate in megabytes at the end of the interval for the host [Sequential = AVG Non-Sequential = SUM]
activewrite	The amount of memory in megabytes actively being written to or by the virtual machine. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]

consumed	The amount of memory in megabytes used on the host. It includes service console memory, VMKernel memory, VMware Infrastructure services memory, and VM memory. [Sequential = AVG Non-Sequential = SUM]
free_mem	The amount of memory in megabytes currently available to be used by the host system [Sequential = LST Non-Sequential = SUM]
heap	The amount of memory in megabytes allocated for the heap [Sequential = AVG Non-Sequential = SUM]
heapfree	The amount of free space in megabytes in the memory heap [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
latency	The percentage of time the virtual machine was waiting to access swapped or compressed memory. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
lowfreethreshold	The threshold of free host physical memory in megabytes. The VMware ESX Server will begin reclaiming memory from virtual machines through ballooning and swapping if the amount of free host physical memory falls below this value. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
overhead	The amount of additional host memory in megabytes allocated to the virtual machine [Sequential = AVG Non-Sequential = SUM]
reservedCapacity	The amount of memory in megabytes reserved by the virtual machines on a VMware host [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shared	The amount of memory in megabytes shared between the virtual machines on a VMware host [Sequential = AVG Non-Sequential = ID]
shared_comm_mem	The total amount of shared common memory in megabytes on the host system [Sequential = LST Non-Sequential = SUM]
size	The amount of memory in megabytes granted [Sequential = AVG Non-Sequential = SUM]
state	The VMKernel threshold for the amount of free memory on the host [Sequential = LST Non-Sequential = SUM]
swapiin	The total amount of memory in megabytes that is swapped in on a VMware host [Sequential = AVG Non-Sequential = SUM]

swpinRate	The rate memory is swapped from the disk into active memory during the collection interval. This value applies to virtual machines. This value is more useful than the swpin statistic to determine if the virtual machine is running slow due to memory swapping, especially when evaluating real-time data. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
swapout	The total amount of memory in megabytes that is swapped out on a VMware host [Sequential = AVG Non-Sequential = SUM]
swapoutRate	The rate memory is swapped from active memory to the disk during the collection interval. This value applies to virtual machines. This value is more useful than the swapout statistic to determine if the virtual machine is running slow due to memory swapping, especially when evaluating real-time data. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
swapped	The amount of memory in megabytes that is used by swap [Sequential = AVG Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
sysUsage	The amount of memory in megabytes used by the VMKernel for core functionality, such as device drivers and other internal usage [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
total_mem	The total amount of physical memory in megabytes on the host system [Sequential = LST Non-Sequential = SUM]
totalCapacity	The total amount of memory reservation in megabytes used by and available for powered-on virtual machines and VMware vSphere services on the host. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
unreserved	The amount of memory in megabytes that is unreserved [Sequential = AVG Non-Sequential = SUM]
vmmemctl	The amount of memory in megabytes used by memory control [Sequential = AVG Non-Sequential = SUM]
zero	The amount of memory in megabytes that is zeroed out [Sequential = AVG Non-Sequential = SUM]

17.1.17. Network Device.by Host System Device Table

The Network Device.by Host System Device table stores data on network device usage by VMware hosts.

Table Field Hierarchy

Class:	Network Device
Subclass:	by Host System Device
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	Network Device.by Host System Device
Open Table Name:	NETDEVBYHOSTSYSDEVIC
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Host Network Device Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
broadcast/s	The total number of broadcast packets received and sent per second by the host system interface. This statistic is available for VMware ESX 5.0.0 and later. Calculated as $\text{broadcast} = \text{broadcastRx} + \text{broadcastTx}$ [Sequential = SUM Non-Sequential = SUM]
broadcastRx/s	The number of broadcast packets received by the host system interface. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
broadcastTx/s	The number of broadcast packets sent by the host system interface. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
dropped	The total number of packets dropped during the collection interval. Calculated as $\text{dropped} = \text{droppedRx} + \text{droppedTx}$ [Sequential = SUM Non-Sequential = SUM]

droppedRx	The number of receive packets dropped during the collection interval. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
droppedTx	The number of sent packets dropped during the collection interval. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KBRx/s	The amount of data received per second in kilobytes (KB) by the host system interface [Sequential = AVG Non-Sequential = SUM]
KBTx/s	The amount of data transmitted per second in kilobytes (KB) by the host system interface [Sequential = AVG Non-Sequential = SUM]
KBx/s	The amount of data transferred per second in kilobytes (KB) by the host system interface [Sequential = AVG Non-Sequential = SUM]
multicast	The total number of multicast packets received and sent during the collection interval. Calculated as $\text{multicast} = \text{multicastRx} + \text{multicastTx}$ [Sequential = SUM Non-Sequential = SUM]
multicastRx	The number of multicast packets received during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
multicastTx	The number of multicast packets sent during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
packets/s	The total number of packets transferred per second by the host system interface [Sequential = AVG Non-Sequential = SUM]
pktsRx/s	The number of packets received per second by the host system interface [Sequential = AVG Non-Sequential = SUM]
pktsTx/s	The number of packets transmitted per second by the host system interface [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the host system interface [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
unknownProtos	The number of frames with unknown protocols received during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]

17.1.18. Network Device.by Virtual Machine Table

The Network Device.by Virtual Machine table stores overall network performance data for VMware virtual machines.

Table Field Hierarchy

Class:	Network Device
Subclass:	by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Network Device.by Virtual Machine
Open Table Name:	NETDEVBYVM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Virtual Machine Network Device Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
broadcast/s	The total number of broadcast packets received and sent per second by the host system interface. This statistic is available for VMware ESX 5.0.0 and later. Calculated as $\text{broadcast} = \text{broadcastRx} + \text{broadcastTx}$ [Sequential = AVG Non-Sequential = SUM]
broadcastRx/s	The number of broadcast packets received per second by the host system interface. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]

broadcastTx/s	The number of broadcast packets sent per second by the host system interface. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
dropped	The total number of packets dropped during the collection interval. Calculated as $\text{dropped} = \text{droppedRx} + \text{droppedTx}$ [Sequential = SUM Non-Sequential = SUM]
droppedRx	The number of receive packets dropped during the collection interval. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
droppedTx	The number of sent packets dropped during the collection interval. This statistic is available for VMware ESX 4.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
errors	The total number of packets with errors received and sent during the collection interval. Calculated as $\text{errors} = \text{errorsRx} + \text{errorsTx}$ [Sequential = SUM Non-Sequential = SUM]
errorsRx	The number of packets with errors received during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
errorsTx	The number of packets with errors sent during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KBRx/s	The amount of data received per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KBTx/s	The amount of data transmitted per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KBx/s	The amount of data transferred per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
multicast	The total number of multicast packets received and sent during the collection interval. Calculated as $\text{multicast} = \text{multicastRx} + \text{multicastTx}$ [Sequential = SUM Non-Sequential = SUM]
multicastRx	The number of multicast packets received during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
multicastTx	The number of multicast packets sent during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]

packets/s	The number of packets transferred per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
pktsRx/s	The number of packets received per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
pktsTx/s	The number of packets transmitted per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the host system interface [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

17.1.19. Network Device.vmnic by Virtual Machine Table

The Network Device.vmnic by Virtual Machine table stores overall network performance data for VMware virtual machines.

Table Field Hierarchy

Class:	Network Device
Subclass:	vmnic by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	Network Device.vmnic by Virtual Machine
Open Table Name:	NETDEVVMNICBYVM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Virtual Machine Network vmnic Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KBRx/s	The amount of data received per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KBTx/s	The amount of data transmitted per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KBx/s	The amount of data transferred per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
packets/s	The number of packets transferred per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
pktsRx/s	The number of packets received per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
pktsTx/s	The number of packets transmitted per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the host system interface [Sequential = ID Non-Sequential = ID]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

17.1.20. Network Device.VMware Summary Table

The Network Device.VMware Summary table stores overall network performance data for VMware hosts.

Table Field Hierarchy

Class:	Network Device
Subclass:	VMware Summary
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	Network Device.VMware Summary
Open Table Name:	NETDEVVMSUM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Host Network Device Summary

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
broadcast/s	The total number of broadcast packets received and sent per second by the host system interface. This statistic is available for VMware ESX 5.0.0 and later. Calculated as $\text{broadcast} = \text{broadcastRx} + \text{broadcastTx}$ [Sequential = AVG Non-Sequential = SUM]

broadcastRx/s	<p>The number of broadcast packets received per second by the host system interface. This statistic is available for VMware ESX 5.0.0 and later.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
broadcastTx/s	<p>The number of broadcast packets sent per second by the host system interface. This statistic is available for VMware ESX 5.0.0 and later.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
dropped	<p>The total number of packets dropped during the collection interval. Calculated as</p> $\text{dropped} = \text{droppedRx} + \text{droppedTx}$ <p>[Sequential = SUM Non-Sequential = SUM]</p>
droppedRx	<p>The number of receive packets dropped during the collection interval. This statistic is available for VMware ESX 4.0.0 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
droppedTx	<p>The number of sent packets dropped during the collection interval. This statistic is available for VMware ESX 4.0.0 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
errors	<p>The total number of packets with errors received and sent during the collection interval. Calculated as</p> $\text{errors} = \text{errorsRx} + \text{errorsTx}$ <p>[Sequential = SUM Non-Sequential = SUM]</p>
errorsRx	<p>The number of packets with errors received during the collection interval. This statistic is available for VMware ESX 5.0.0 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
errorsTx	<p>The number of packets with errors sent during the collection interval. This statistic is available for VMware ESX 5.0.0 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
Interval	<p>The expected sampling interval in seconds</p> <p>[Sequential = SUM Non-Sequential = AVG]</p>
KBRx/s	<p>The amount of data received per second in kilobytes (KB)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
KBTx/s	<p>The amount of data transmitted per second in kilobytes (KB)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
KBx/s	<p>The amount of data transferred per second in kilobytes (KB)</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
multicast	<p>The total number of multicast packets received and sent during the collection interval. Calculated as</p> $\text{multicast} = \text{multicastRx} + \text{multicastTx}$ <p>[Sequential = SUM Non-Sequential = SUM]</p>
multicastRx	<p>The number of multicast packets received during the collection interval. This statistic is available for VMware ESX 5.0.0 and later.</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>

multicastTx	The number of multicast packets sent during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]
packets/s	The total number of packets transferred per second [Sequential = AVG Non-Sequential = SUM]
pktsRx/s	The total number of packets received per second [Sequential = AVG Non-Sequential = SUM]
pktsTx/s	The total number of packets transmitted per second [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
unknownProtos	The number of frames with unknown protocols received during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = SUM Non-Sequential = SUM]

17.1.21. VMware.Availability by Virtual Machine Table

The VMware.Availability by Virtual Machine table stores high level status data.

Table Field Hierarchy

Class:	VMware
Subclass:	Availability by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	VMware.Availability by Virtual Machine
Open Table Name:	VMAVAILABILITYBYVM
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Virtual Machine Availability

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
heartbeat	The number of heartbeats in the collection period. The heartbeat represents the overall health of the guest operating system. [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
uptime_t	The total time in days elapsed since the last virtual machine reboot [Sequential = LST Non-Sequential = SUM]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

17.1.22. VMware.Host Configuration Table

The VMware.Host Configuration table stores configuration data for VMware hosts.

Table Field Hierarchy

Class:	VMware
Subclass:	Host Configuration
IT Resource Name:	/TeamQuest/System/VMware/vCenter/ <i>vCenter</i>
TeamQuest Table Name:	VMware.Host Configuration
Open Table Name:	VMHOSTCONF
Collection interval:	N/A
Default retention:	1 year
Table type:	State
Derived tables using fields from this table:	Cluster CPU Summary Cluster Memory Summary Cluster Resource Allocation Cluster Virtual Machine Operations Host Block Device Summary Host Block Device Usage Host CPU Summary Host CPU Usage Host Memory Summary Host Network Device Summary Host Network Device Usage Virtual Machine Availability Virtual Machine Block Device Usage Virtual Machine Configuration Virtual Machine CPU Usage Virtual Machine Memory Usage Virtual Machine Network Device Usage Virtual Machine Network vmnic Usage Virtual Machine Virtual CPU Usage

Statistic Name	Description
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Non-Sequential = ID]

directoryPath	<p>The path of the directory under the VMware vCenter where the VMware cluster resides. This path name consists of zero or more administrator-defined folder names separated by slashes, followed by the datacenter name. The datacenter name can be followed by zero or more administrator-defined folder names separated by slashes. Often, the directoryPath is the datacenter name. This is the path name that was configured at the time the record was stored.</p> <p>[Non-Sequential = ID]</p>
Host	<p>The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters is truncated.</p> <p>[Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Non-Sequential = ID]</p>
vCenter	<p>The name of the VMware vCenter server</p> <p>[Non-Sequential = ID]</p>

17.1.23. VMware.Host Status Table

The VMware.Host Status table stores performance data for VMware datastores assigned to individual hosts.

Table Field Hierarchy

Class:	VMware
Subclass:	Host Status
IT Resource Name:	/TeamQuest/System/VMware/vCenter/vCenter
TeamQuest Table Name:	VMware.Host Status
Open Table Name:	VMHOSTSTATUS
Collection interval:	Based on the collection period
Default retentions:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
collectionAvailable	The percentage of time that data collection was possible with the VMware Infrastructure Agent. To accurately determine this value, the vCenter Host configuration Collection Method field must be set to Through vCenter in the TeamQuest Manager interface. Collection is available if the following conditions occur: connected is 100.00 poweredOn is 100.00 [Sequential = AVG Non-Sequential = AVG]
connected	The percentage of time the VMware host was connected to the VMware vCenter server, sampled at the end of each collection interval [Sequential = AVG Non-Sequential = AVG]
disconnected	The percentage of time the VMware host was disconnected from the VMware vCenter server, sampled at the end of each collection interval [Sequential = AVG Non-Sequential = AVG]
Host	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
maintenanceMode	The percentage of time the VMware host was in maintenance mode at the end of each collection interval [Sequential = AVG Non-Sequential = AVG]

VMware Systems

notResponding	The percentage of time the VMware host was not responding to the VMware vCenter server, sampled at the end of each collection interval [Sequential = AVG Non-Sequential = AVG]
poweredOff	The percentage of time the VMware host was powered off, sampled at the end of each collection interval [Sequential = AVG Non-Sequential = AVG]
poweredOn	The percentage of time the VMware host was powered on, sampled at the end of each collection interval [Sequential = AVG Non-Sequential = AVG]
powerUnknown	The percentage of time the VMware host was in an unknown state at the end of each collection interval [Sequential = AVG Non-Sequential = AVG]
standby	The percentage of time the VMware host was on standby, sampled at the end of each collection interval [Sequential = AVG Non-Sequential = AVG]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.1.24. VMware.Storage Configuration Table

The VMware.Storage Configuration table stores configuration data for VMware datastore storage.

Table Field Hierarchy

Class:	VMware
Subclass:	Storage Configuration
IT Resource Name:	/TeamQuest/System/VMware/vCenter/ <i>vCenter</i>
TeamQuest Table Name:	VMware.Storage Configuration
Open Table Name:	VMSTORAGECONF
Collection interval:	N/A
Default retentions:	1 year
Table type:	State
Derived tables using fields from this table:	Storage Datastore by Host System Storage Datastore Summary Storage Datastore by Virtual Machine Datastore Cluster File Type Usage Datastore Cluster Summary Datastore Cluster Usage by Virtual Machine

Statistic Name	Description
Datacenter	The name of the datacenter to which the datastore belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]
Datastore_Cluster	The name of the VMware datastore cluster to which the data applies [Sequential = ID Non-Sequential = ID]
Datastore_ClusterId	The identifier assigned to the VMware datastore cluster that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.1.25. VMware.Support Metrics Table

The VMware.Support Metrics table contains data used by customer support when troubleshooting customer issues.

Table Field Hierarchy

Class:	VMware
Subclass:	Support Metrics
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	VMware.Support Metrics
Open Table Name:	VMSUPPORTMETRICS
Collection interval:	Based on the collection period
Default retentions:	1 week at collection period interval
Table type:	Performance
Derived tables using fields from this table:	N/A

17.1.26. VMware.Virtual_Machines Table

The VMware.Virtual_Machines table stores configuration data about virtual machines.

Table Field Hierarchy

Class:	VMware
Subclass:	Virtual_Machines
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	VMware.Virtual_Machines
Open Table Name:	VMVMS
Collection interval:	N/A
Default retention:	1 year
Table type:	State
Derived tables using fields from this table:	Virtual Machine Configuration

Statistic Name	Description
Configuration_File	The configuration file for the virtual machine [Non-Sequential = ID]
CPU_Limit	The cap on the consumption of CPU time by the virtual machine, measured in megahertz (MHz). A value of zero indicates no limit on CPU consumption. [Non-Sequential = SUM]
CPU_Reservation	The number of CPU cycles reserved for the virtual machine, measured in megahertz (MHz) [Non-Sequential = SUM]
CPU_Shares	The CPU share allocation for the virtual machine [Non-Sequential = SUM]
Disk_Shares	The disk share allocation for the virtual machine [Non-Sequential = SUM]
ESX_Server	The version of VMware ESX Server [Non-Sequential = ID]
htSharing	Specifies how the VCPUs of a virtual machine are allowed to share physical cores on a hyperthreaded system. Values can be any, internal, or none. [Non-Sequential = ID]
Memory_Limit	The cap on the memory consumption by this virtual machine, measured in megabytes. A value of zero indicates no fixed limit on memory consumption. [Non-Sequential = SUM]
Memory_Reservation	The amount of memory reserved for the virtual machine, measured in megabytes [Non-Sequential = SUM]

VMware Systems

Memory_Shares	The memory share allocation for the virtual machine [Non-Sequential = SUM]
OS	The virtual machine operating system name [Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
VM	The virtual machine name [Non-Sequential = ID]
VMID	The virtual machine identifier [Non-Sequential = ID]
VCPU_Count	The number of virtual processors for the virtual machine [Non-Sequential = SUM]

17.1.27. VMware Cluster.CPU Summary Table

The VMware Cluster.CPU Summary table stores CPU usage data by VMware clusters.

The historical statistic interval is the “Past-day” statistic interval configured in the VMware vCenter server. For example, if the “Past-day” statistic interval in the VMware vCenter server is configured as a 5-minute interval, the historical statistic interval is 3 days at 5-minute intervals.

Table Field Hierarchy

Class:	VMware Cluster
Subclass:	CPU Summary
IT Resource Name:	/TeamQuest/System/VMware/Cluster/ <i>Cluster(ClusterId)</i>
TeamQuest Table Name:	VMware Cluster.CPU Summary
Open Table Name:	VMCLSTRCPUSUM
Collection interval:	Based on the historical statistic interval
Default retention:	3 days at historical statistic interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Cluster CPU Summary

Statistic Name	Description
%busy	The percentage of elapsed CPU time the processors were busy across all of the virtual machines in the VMware cluster. A value of <N/A> is displayed when the VMware Distributed Resource Scheduler (DRS) is disabled for the VMware cluster. [Sequential = AVG Non-Sequential = AVG]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the data applies [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a VMware cluster is renamed. [Sequential = ID Non-Sequential = ID]

effectivecpu	<p>The total amount of available CPU in megahertz (MHz) of all of the hosts within a VMware cluster. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. Calculated as</p> $\text{effectivecpu} = \text{aggregate host CPU capacity} - (\text{VMKernel} + \text{service console CPU} + \text{other service CPU})$ <p>[Sequential = AVG Non-Sequential = AVG]</p>
Interval	<p>The expected sampling interval in seconds</p> <p>[Sequential = SUM Non-Sequential = AVG]</p>
reservedCapacity	<p>The total CPU capacity in megahertz (MHz) reserved by all of the virtual machines within a VMware cluster. This statistic is only available for VMware vCenter server statistic level 2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
total(MHz)	<p>The total amount of CPU in megahertz (MHz) of all of the hosts within the VMware cluster. The maximum value is equal to the frequency of the processors multiplied by the number of cores. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
usage(MHz)	<p>The CPU usage in megahertz (MHz) of all of the powered-on virtual machines in a VMware cluster. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
vCenter	<p>The name of the VMware vCenter server</p> <p>[Sequential = ID Non-Sequential = ID]</p>

17.1.28. VMware Cluster.Memory Summary Table

The VMware Cluster.Memory summary table stores memory usage data by VMware clusters.

The historical statistic interval is the “Past-day” statistic interval configured in the VMware vCenter server. For example, if the “Past-day” statistic interval in the VMware vCenter server is configured as a 5-minute interval, the historical statistic interval is 3 days at 5-minute intervals.

Table Field Hierarchy

Class:	VMware Cluster
Subclass:	Memory Summary
IT Resource Name:	/TeamQuest/System/VMware/Cluster/ <i>Cluster(ClusterId)</i>
TeamQuest Table Name:	VMware Cluster.Memory Summary
Open Table Name:	VMCLSTRMEMSUM
Collection interval:	Based on the historical statistic interval
Default retention:	3 days at historical statistic interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Cluster Memory Summary

Statistic Name	Description
%usage	The percentage of total available memory that is used. A value of <N/A> is displayed when the VMware Distributed Resource Scheduler (DRS) is disabled for the VMware cluster. Calculated as $\text{memory usage} = (\text{memory consumed} + \text{memory overhead}) / \text{effectivemem}$ [Sequential = AVG Non-Sequential = AVG]
active	The working set size estimate in megabytes at the end of the interval for the VMware cluster. This statistic is only available for VMware vCenter server statistic level 2. [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]

balloon	<p>The amount of memory in megabytes allocated by the virtual machine memory control driver (vmmemctl), which is installed with VMware Tools. The vmmemctl is a memory management driver that controls ballooning. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
consumed	<p>The amount of host memory in megabytes consumed by all of the powered-on virtual machines for guest memory within a VMware cluster</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Cluster	<p>The name of the VMware cluster to which the data applies</p> <p>[Sequential = ID Non-Sequential = ID]</p>
ClusterId	<p>The unique identifier for a VMware cluster provided by the VMware API. This identifier does not change when a cluster is renamed.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
effectivemem	<p>The total amount of memory in megabytes of all of the hosts within a VMware cluster that are available for the virtual machine memory and the virtual machine overhead memory. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = AVG]</p>
granted	<p>The amount of all granted memory in megabytes for all of the powered-on virtual machines. This statistic is only available for VMware vCenter server statistic level 2.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Interval	<p>The expected sampling interval in seconds</p> <p>[Sequential = SUM Non-Sequential = AVG]</p>
overhead	<p>The amount of additional host memory in megabytes allocated to the virtual machine. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
reservedCapacity	<p>The amount of memory in megabytes reserved by the virtual machines on a VMware host. This statistic is only available for VMware vCenter server statistic level 2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
shared	<p>The amount of guest physical memory in megabytes shared with other virtual machines. This value includes the amount of zero memory.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
swapused	<p>The amount of memory in megabytes that is used by swap. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. This statistic is only available for VMware vCenter server statistic level 2.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>

sysUsage	<p>The amount of memory in megabytes used by the VMKernel for core functionality, such as device drivers and other internal usage components. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. This statistic is only available for VMware vCenter server statistic level 2 and for VMware vCenter Server 4.0 and lower.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
total	<p>The total amount of memory in megabytes of all of the hosts within a VMware cluster that are available for the virtual machine memory and virtual machine overhead memory. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
vCenter	<p>The name of the VMware vCenter server</p> <p>[Sequential = ID Non-Sequential = ID]</p>
zero	<p>The amount of memory in megabytes that is zeroed out. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>

17.1.29. VMware Cluster.Resource Allocation Table

The VMware Cluster.Resource Allocation table stores usage of VMware host resources by VMware clusters.

The historical statistic interval is the “Past-day” statistic interval configured in the VMware vCenter server. For example, if the “Past-day” statistic interval in the VMware vCenter server is configured as a 5-minute interval, the historical statistic interval is 3 days at 5-minute intervals.

Table Field Hierarchy

Class:	VMware Cluster
Subclass:	Resource Allocation
IT Resource Name:	/TeamQuest/System/VMware/Cluster/ <i>Cluster(ClusterId)</i>
TeamQuest Table Name:	VMware Cluster.Resource Allocation
Open Table Name:	VMCLSTRRESALLOC
Collection interval:	Based on the historical statistic interval
Default retention:	3 days at historical statistic interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Cluster Resource Allocation

Statistic Name	Description
activeHosts	The number of active hosts in the VMware cluster [Sequential = LST Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the data applies [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. [Sequential = ID Non-Sequential = ID]
cpuAvailable	The total amount of CPU in megahertz (MHz) available to satisfy a reservation for all of the virtual machines and resource pools in the cluster [Sequential = LST Non-Sequential = SUM]
cpuReserved	The total amount of CPU in megahertz (MHz) that has been used to satisfy the reservation requirements of all of the descendants of the virtual machines and resource pools in the cluster [Sequential = LST Non-Sequential = SUM]

Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
memAvailable	The total amount of memory in megabytes that is available to satisfy the reservation for all of the virtual machines and resource pools in the cluster [Sequential = LST Non-Sequential = SUM]
memOverhead	The total amount of memory in megabytes that has been used to satisfy the reservation requirements of all of the descendants of the running virtual machines in the cluster [Sequential = LST Non-Sequential = SUM]
memReserved	The total amount of memory in megabytes that has been used to satisfy the reservation requirements of all of the descendants of the virtual machines and resource pools in the cluster [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totalCpu	The total available CPU in megahertz (MHz) of all of the hosts within the cluster [Sequential = LST Non-Sequential = SUM]
totalMem	The total amount of memory in megabytes of all of the hosts within the cluster that is available for use for virtual machine memory and virtual machine overhead memory [Sequential = LST Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.1.30. VMware Cluster.Virtual Machine Operations

The VMware Cluster.Virtual Machine Operations table stores data on the number and type of operations performed on the virtual machine.

The historical statistic interval is the “Past-day” statistic interval configured in the VMware vCenter server. For example, if the “Past-day” statistic interval in the VMware vCenter server is configured as a 5-minute interval, the historical statistic interval is 3 days at 5-minute intervals.

Table Field Hierarchy

Class:	VMware Cluster
Subclass:	Virtual Machine Operations
IT Resource Name:	/TeamQuest/System/VMware/Cluster/ <i>Cluster(ClusterId)</i>
TeamQuest Table Name:	VMware Cluster.Virtual Machine Operations
Open Table Name:	VMCLSTRVMOPERATIONS
Collection interval:	Based on the historical statistic interval
Default retention:	3 days at historical statistic interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Cluster Virtual Machine Operations

Statistic Name	Description
activeVMs	The number of active virtual machines in the VMware cluster [Sequential = LST Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
clone	The number of virtual machine clone operations [Sequential = SUM Non-Sequential = SUM]
Cluster	The name of the cluster to which the data applies [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. [Sequential = ID Non-Sequential = ID]
create	The number of virtual machine create operations [Sequential = SUM Non-Sequential = SUM]

datastoreChange	The number of datastore change operations for powered-off and suspended virtual machines [Sequential = SUM Non-Sequential = SUM]
delete	The number of virtual machine delete operations [Sequential = SUM Non-Sequential = SUM]
failover	The number of virtual machine failover operations. This statistic is only available for clusters with VMware High Availability (HA) on. [Sequential = LST Non-Sequential = SUM]
guestReboot	The number of virtual machine guest reboot operations [Sequential = SUM Non-Sequential = SUM]
guestShutdown	The number of virtual machine guest shutdown operations [Sequential = SUM Non-Sequential = SUM]
hostChange	The total number of host change operations for powered-down and suspended virtual machines [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
powerOff	The number of virtual machine power-off operations [Sequential = SUM Non-Sequential = SUM]
powerOn	The number of virtual machine power-on operations [Sequential = SUM Non-Sequential = SUM]
reconfigure	The number of virtual machine reconfigure operations [Sequential = SUM Non-Sequential = SUM]
register	The number of virtual machine register operations [Sequential = SUM Non-Sequential = SUM]
reset	The number of virtual machine reset operations [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
standByGuest	The number of virtual machine standby guest operations [Sequential = SUM Non-Sequential = SUM]
storageMotion	The number of migrations with storage vMotion or datastore change operations for all of the powered-on virtual machines [Sequential = SUM Non-Sequential = SUM]
suspend	The number of virtual machine suspend operations [Sequential = SUM Non-Sequential = SUM]
templateDeploy	The number of virtual machine template deploy operations [Sequential = SUM Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
unregister	The number of virtual machine unregister operations [Sequential = SUM Non-Sequential = SUM]

vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
vMotion	The number of migrations with vMotion for powered-on virtual machines [Sequential = SUM Non-Sequential = SUM]

17.1.31. VMware Datastore.File Type Usage by Datacenter Table

The VMware Datastore.File Type Usage by Datacenter table stores data on how datastore space is used by datastore and file type.

Note: The minimum collection period for VMware datastores is 5 minutes.

Table Field Hierarchy

Class:	VMware Datastore
Subclass:	File Type Usage by Datacenter
IT Resource Name:	/TeamQuest/System/VMware/Datastore/Datastore
TeamQuest Table Name:	VMware Datastore.File Type Usage by Datacenter
Open Table Name:	VMDSFILETYPEUSAGEBYD
Collection interval:	Based on the collection period
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Datastore Cluster File Type Usage

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
capacity	The configured size of the VMware datastore in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the datastore belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]

Datastore_Cluster	The name of the VMware datastore cluster to which the data applies [Sequential = ID Non-Sequential = ID]
Datastore_ClusterId	The identifier assigned to the VMware datastore cluster that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
free	The amount of free physical VMware datastore space in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
other	The amount of disk space used by all other non-virtual machine files in gigabytes (GB), such as documentation files and backup files [Sequential = LST Non-Sequential = SUM]
other_vm_files	The amount of disk space used by all other virtual machine files in gigabytes (GB), such as configuration files and log files [Sequential = LST Non-Sequential = SUM]
provisioned	The amount of physical space in gigabytes (GB) provisioned by an administrator for the VMware datastore. This value is the maximum storage size to which files on the VMware datastore can grow. [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
swapfiles	The amount of disk space in gigabytes (GB) used by swap files. Swap files are used to back up the virtual machine physical memory. [Sequential = LST Non-Sequential = SUM]
snapshots	The amount of disk space in gigabytes (GB) used by virtual machine snapshot files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. The virtual machine snapshot files have the file extension of .vmsn. [Sequential = LST Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]

vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
virtual_disks	The amount of disk space in gigabytes (GB) used by virtual disk files and delta disk files. Virtual disk files store the contents of the virtual machines hard disk drive, including information that is written to the virtual machines hard disk (the operating system, program files, and data files). The virtual disk files have the file extension of .vmdk and appear as a physical disk drive on a guest operating system. Delta disk files store the updates made by the virtual machine to the virtual disks after a snapshot is taken. [Sequential = LST Non-Sequential = SUM]

17.1.32. VMware Datastore.Summary Table

The VMware Datastore.Summary table stores datastore usage data by datastore.

Note: The minimum collection period for VMware datastores is 5 minutes.

Table Field Hierarchy

Class:	VMware Datastore
Subclass:	Summary
IT Resource Name:	/TeamQuest/System/VMware/Datastore/Datastore
TeamQuest Table Name:	VMware Datastore.Summary
Open Table Name:	VMDSSUM
Collection interval:	Based on the collection period
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Datastore Cluster Summary

Statistic Name	Description
%free	The percentage of free physical VMware datastore space [Sequential = LST Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
capacity	The configured size of the VMware datastore in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]

Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]
free	The amount of free physical VMware datastore space in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
used	The amount of the physical VMware datastore space in use in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.1.33. VMware Datastore.Usage by Virtual Machine Table

The VMware Datastore.Usage by Virtual Machine table stores datastore usage by virtual machines.

Note: *The minimum collection period for VMware datastores is 5 minutes.*

Table Field Hierarchy

Class:	VMware Datastore
Subclass:	Usage by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Datastore/Datastore
TeamQuest Table Name:	VMware Datastore.Usage by Virtual Machine
Open Table Name:	VMDSUSAGEBYVM
Collection interval:	Based on the collection period
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Datastore Cluster Usage by Virtual Machine

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Datacenter	The name of the datacenter to which the datastore belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]
Host	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
provisioned	The amount of storage space in gigabytes (GB) set aside for use by a virtual machine [Sequential = LST Non-Sequential = SUM]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
unshared	The amount of storage space in gigabytes (GB) associated exclusively with a virtual machine [Sequential = LST Non-Sequential = SUM]
used	The amount of the target VMware datastore space in gigabytes (GB) used by a virtual machine [Sequential = LST Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

17.1.34. VMware Resource Pool.CPU Summary Table

The VMware Resource Pool.CPU Summary table stores CPU usage data by VMware resource pools.

The historical statistic interval is the “Past-day” statistic interval configured in the VMware vCenter server. For example, if the “Past-day” statistic interval in the VMware vCenter server is configured as a 5-minute interval, the historical statistic interval is 3 days at 5-minute intervals.

Table Field Hierarchy

Class:	VMware Resource Pool
Subclass:	CPU Summary
IT Resource Name:	/TeamQuest/System/VMware/ResourcePool/ <i>Resource_pool(Resource Pool_Id)</i>
TeamQuest Table Name:	VMware Resource Pool.CPU Summary
Open Table Name:	VMRESPOOLCPUSUM
Collection interval:	Based on the historical statistic interval
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Resource Pool CPU Summary

Statistic Name	Description
%busy	The percentage of elapsed CPU time the processors were busy across all of the virtual machines in the resource pool [Sequential = AVG Non-Sequential = AVG]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
capacityContention	The percentage of time the virtual machine is unable to run because it is contending for access to the physical CPUs. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
capacityDemand	The amount of CPU resources a virtual machine would use if there were no CPU contention or CPU limit. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]

capacityEntitlement	The amount of CPU resources devoted by the VMware ESX scheduler to virtual machines and resource pools. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
Cluster	The name of the cluster to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware host. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for the cluster provided by the VMware API. This identifier does not change when a cluster is renamed. [Sequential = ID Non-Sequential = ID]
corecountContention	The amount of time the virtual machine was ready to run but was unable to run due to co-scheduling constraints. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
corecountProvisioned	The number of virtual processors or physical cores provisioned to the entity. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the resource pool belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Host	The name of the host system to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware cluster. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]
ResourcePool	The name of the resource pool to which the data applies [Sequential = ID Non-Sequential = ID]
ResourcePool_Id	The unique identifier for a resource pool provided by the VMware API. This identifier does not change when a resource pool is renamed. [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
usage(MHz)	The CPU usage in megahertz (MHz) over the collection interval [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.1.35. VMware Resource Pool.Memory Summary Table

The VMware Resource Pool.Memory Summary table stores memory usage data by VMware resource pools.

The historical statistic interval is the “Past-day” statistic interval configured in the VMware vCenter server. For example, if the “Past-day” statistic interval in the VMware vCenter server is configured as a 5-minute interval, the historical statistic interval is 3 days at 5-minute intervals.

Table Field Hierarchy

Class:	VMware Resource Pool
Subclass:	Memory Summary
IT Resource Name:	/TeamQuest/System/VMware/ResourcePool/ <i>Resource_Pool_Id</i>
TeamQuest Table Name:	VMware Resource Pool.Memory Summary
Open Table Name:	VMRESPOOLMEMSUM
Collection interval:	Based on the historical statistic interval
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Resource Pool Memory Summary

Statistic Name	Description
%usage	The percentage of memory usage over the collection interval. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = AVG]
active	The working set size estimate in megabytes at the end of the interval for the resource pool [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
balloon	The amount of memory in megabytes allocated by the virtual machine memory control driver (vmmemctl), which is installed with VMware Tools. The vmmemctl is a memory management driver that controls ballooning. [Sequential = AVG Non-Sequential = SUM]

capacityContention	The percentage of time the virtual machine is waiting to access swapped or compressed memory. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
capacityEntitlement	The amount of host physical memory devoted by the VMware ESX scheduler to the virtual machine. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
capacityUsage	The amount of physical memory in megabytes used by the virtual machine. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
Cluster	The name of the cluster to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware host. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for the cluster provided by the VMware API. This identifier does not change when a cluster is renamed. [Sequential = ID Non-Sequential = ID]
compressed	The amount of memory in megabytes compressed by the VMware ESX Server. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
compressionRate	The rate of memory compression for the virtual machine. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
consumed	The amount of memory in megabytes consumed by a resource pool, host, or virtual machine [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the resource pool belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
decompressionRate	The rate of memory decompression for the virtual machine. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
granted	The amount of all granted memory in megabytes for all of the powered-on virtual machines. This statistic is only available for VMware vCenter server statistic level 2. [Sequential = AVG Non-Sequential = SUM]
Host	The name of the host system to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware cluster. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]

overhead	The amount of additional resource pool memory in megabytes allocated to the virtual machine [Sequential = AVG Non-Sequential = SUM]
ResourcePool	The name of the resource pool to which the data applies [Sequential = ID Non-Sequential = ID]
ResourcePool_Id	The unique identifier for a resource pool provided by the VMware API. This identifier does not change when a resource pool is renamed. [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shared	The amount of virtual machine memory in megabytes that is shared with other virtual machines, relative to a single virtual machine or to all powered-on virtual machines on a host [Sequential = AVG Non-Sequential = SUM]
swapped	The amount of memory in megabytes currently swapped to the VMware File System 3 (VMFS3) swap file [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
zero	The amount of memory in megabytes that is zeroed out. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]

17.1.36. VMware Resource Pool.Resource Allocation Table

The VMware Resource Pool.Resource Allocation table stores usage of VMware resource pools.

The historical statistic interval is the “Past-day” statistic interval configured in the VMware vCenter server. For example, if the “Past-day” statistic interval in the VMware vCenter server is configured as a 5-minute interval, the historical statistic interval is 3 days at 5-minute intervals.

Table Field Hierarchy

Class:	VMware Resource Pool
Subclass:	Resource Allocation
IT Resource Name:	/TeamQuest/System/VMware/ResourcePool/ <i>Resource_Pool_Id</i>
TeamQuest Table Name:	VMware Resource Pool.Resource Allocation
Open Table Name:	VMRESPOOLRESALLOC
Collection interval:	Based on the historical statistic interval
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Resource Pool Resource Allocation

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
activeVMs	The number of active virtual machines in a resource pool [Sequential = LST Non-Sequential = SUM]
childResourcePool	The number of child resource pools in a resource pool [Sequential = LST Non-Sequential = SUM]
Cluster	The name of the cluster to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware host. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for the cluster provided by the VMware API. This identifier does not change when a cluster is renamed. [Sequential = ID Non-Sequential = ID]
cpuAvailable	The total amount of CPU in megahertz (MHz) that is available to satisfy a reservation requirement for all virtual machines in a resource pool [Sequential = LST Non-Sequential = LST]

cpuReserved	The total amount of CPU in megahertz (MHz) that is used to satisfy the reservation requirements of all descendants of the resource pools and virtual machines in a resource pool [Sequential = LST Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the resource pool belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Host	The name of the host system to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware cluster. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]
memAvailable	The total amount of memory in megabytes that is available to satisfy the reservation requirements for all virtual machines in a resource pool [Sequential = LST Non-Sequential = LST]
memOverhead	The total amount of memory in megabytes used to satisfy the reservation requirements of all descendants of the virtual machines in a resource pool or any of the child resource pools [Sequential = LST Non-Sequential = SUM]
memReserved	The total amount of memory in megabytes used to satisfy the reservation requirements of all descendants of the resource pools and virtual machines in a resource pool [Sequential = LST Non-Sequential = SUM]
ResourcePool	The name of the resource pool to which the data applies [Sequential = ID Non-Sequential = ID]
ResourcePool_Id	The unique identifier for a resource pool provided by the VMware API. This identifier does not change when a resource pool is renamed. [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.1.37. VMware Storage.Adapter by Host System Table

The VMware Storage.Adapter by Host System table stores storage adapter I/O operation data by host.

Table Field Hierarchy

Class:	VMware Storage
Subclass:	Adapter by Host System
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	VMware Storage.Adapter by Host System
Open Table Name:	VMSTORADAPTERBYHOST
Collection interval:	Based on the collection period
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Storage Adapter by Host System

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgQueueLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per command from the perspective of a storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of a storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of a storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
cmds/s	The number of commands (requests) issued per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]

KB_read/s	The amount of data read per second in kilobytes (KB) by the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data transferred per second in kilobytes (KB) by the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
oIOsPct	The percentage of I/O operations that have been issued but have not yet completed. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
outstandingIOs	The number of I/O operations that have been issued but have not yet completed. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
queued	The number of I/O operations waiting to be issued. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
queueDepth	The maximum number of I/O operations that can be outstanding at a time. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
reads/s	The number of read requests issued per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Storage_Adapter	The name of the storage adapter [Sequential = ID Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
throughputContention	The average amount of time in milliseconds for an I/O operation to complete. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = AVG]
throughputUsage	The average amount of data transferred per second in kilobytes by the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
writes/s	The total number of write requests issued per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]

17.1.38. VMware Storage.Adapter Summary Table

The VMware Storage.Adapter Summary table stores I/O operation data summarized by storage adapters.

Table Field Hierarchy

Class:	VMware Storage
Subclass:	Adapter Summary
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	VMware Storage.Adapter Summary
Open Table Name:	VMSTORADAPTERSUM
Collection interval:	Based on the collection period
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Storage Adapter Summary

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
highestLatency	The highest latency in milliseconds of the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = MAX Non-Sequential = MAX]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

17.1.39. VMware Storage.Datastore by Host System

The VMware Storage.Datastore by Host System table stores I/O information for the datastore from the perspective of the VMware host.

Table Field Hierarchy

Class:	VMware Storage
Subclass:	Datastore by Host System
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	VMware Storage.Datastore by Host System
Open Table Name:	VMSTORDSBYHOST
Collection interval:	Based on the collection period
Default retention:	Based on the collection period
Table type:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Derived tables using fields from this table:	Storage Datastore by Host System

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the datastore [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the datastore [Sequential = AVG Non-Sequential = AVG]
cmds	The number of commands (requests) issued per second to the datastore. Calculated as cmds= reads + writes [Sequential = SUM Non-Sequential = SUM]

cmds/s	<p>The number of commands (requests) issued per second to the datastore. Calculated as</p> $\text{cmds/s} = \text{reads/s} + \text{writes/s}$ <p>[Sequential = SUM Non-Sequential = SUM]</p>
datastoreIops	<p>The aggregated number of storage I/O operations on the datastore</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
datastoreMaxQueueDepth	<p>The maximum number of storage I/O operations supported by the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreNormalReadLatency	<p>The storage DRS normalized read latency for the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreNormalWriteLatency	<p>The storage DRS normalized write latency for the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreReadBytes	<p>The number of storage DRS bytes read by the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreReadIops	<p>The number of storage DRS read I/O operations of the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreReadLoadMetric	<p>The storage DRS read workload metric of the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreReadOIO	<p>The number of outstanding read requests by the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreWriteBytes	<p>The number of storage DRS bytes written to the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreWriteIops	<p>The number of storage DRS write I/O operations of the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreWriteLoadMetric	<p>The storage DRS write workload metric of the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
datastoreWriteOIO	<p>The number of outstanding write requests by the datastore</p> <p>[Sequential = LST Non-Sequential = LST]</p>
highestLatency	<p>The highest latency value across all datastores used by the host</p> <p>[Sequential = MAX Non-Sequential = MAX]</p>
Interval	<p>The expected sampling interval in seconds</p> <p>[Sequential = SUM Non-Sequential = AVG]</p>
KB_read/s	<p>The amount of data read per second in kilobytes (KB) by the datastore</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
KB_write/s	<p>The amount of data written per second in kilobytes (KB) by the datastore</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
KB/s	<p>The amount of data requests per second in kilobytes (KB) by the datastore. Calculated as</p> $\text{KB/s} = \text{KB_reads} + \text{KB_writes}$ <p>[Sequential = AVG Non-Sequential = SUM]</p>

reads	The number of read requests issued to the datastore during the collection interval [Sequential = SUM Non-Sequential = SUM]
reads/s	The number of read requests issued per second to the datastore [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
sizeNormalized DatastoreLatency	The normalized latency in milliseconds on the datastore. Data for all virtual machines is combined into this statistic. [Sequential = AVG Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to the datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
writes	The number of write requests issued per second to the datastore [Sequential = SUM Non-Sequential = SUM]
writes/s	The number of write requests issued to the datastore during the collection interval [Sequential = AVG Non-Sequential = SUM]

17.1.40. VMware Storage.Datastore by Virtual Machine

The VMware Storage.Datastore by Virtual Machine table stores I/O information for the datastore from the perspective of the virtual machine.

Table Field Hierarchy

Class:	VMware Storage
Subclass:	Datastore by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
TeamQuest Table Name:	VMware Storage.Datastore by Virtual Machine
Open Table Name:	VMSTORDSBYVM
Collection interval:	Based on the collection period
Default retention:	Based on the collection period
Table type:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Derived tables using fields from this table:	Storage Datastore by Virtual Machine

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgLatency	The average amount of time in milliseconds taken to complete a command request (queue and disk service time) by the host system disk [Sequential = AVG Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the datastore [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the datastore [Sequential = AVG Non-Sequential = AVG]
cmds	The number of commands (requests) issued per second to the datastore. Calculated as cmds= reads + writes [Sequential = SUM Non-Sequential = SUM]
cmds/s	The number of commands (requests) issued per second to the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]

highestLatency	The highest latency value across all datastores used by the host [Sequential = MAX Non-Sequential = MAX]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data requests per second in kilobytes (KB) by the datastore. Calculated as $KB/s = KB_reads + KB_writes$ [Sequential = AVG Non-Sequential = SUM]
reads	The number of read requests issued to the datastore during the collection interval [Sequential = SUM Non-Sequential = SUM]
reads/s	The number of read requests issued per second to the datastore [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to the datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
writes	The number of write requests issued per second to the datastore [Sequential = SUM Non-Sequential = SUM]
writes/s	The number of write requests issued to the datastore during the collection interval [Sequential = AVG Non-Sequential = SUM]

17.1.41. VMware Storage.Datastore Summary

The VMware Storage.Datastore Summary table stores I/O information for the datastore.

Table Field Hierarchy

Class:	VMware Storage
Subclass:	Datastore Summary
IT Resource Name:	/TeamQuest/System/VMware/Datastore/ <i>Datastore</i>
TeamQuest Table Name:	VMware Storage.Datastore Summary
Open Table Name:	VMSTORDSBYDS
Collection interval:	Based on the collection period
Default retention:	Based on the collection period
Table type:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Derived tables using fields from this table:	Storage Datastore Summary

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data requests per second in kilobytes (KB) by the datastore. Calculated as KB/s= KB_reads + KB_writes [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to the datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.1.42. VMware Storage.Path by Host System Table

The VMware Storage.Path by Host System table stores storage path I/O operation data by host.

Table Field Hierarchy

Class:	VMware Storage
Subclass:	Path by Host System
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	VMware Storage.Path by Host System
Open Table Name:	VMSTORPATHBYHOST
Collection interval:	Based on the collection period
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Storage Path by Host System

Statistic Name

Description

Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
busRst/s	The number of SCSI bus resets per second that occurred on the storage path. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
cmds/s	The number of commands (requests) issued per second to the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
cmdsAbrt/s	The number of SCSI commands aborted by the storage path per second. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]

Device_Name	The name of the host system device [Sequential = ID Non-Sequential = ID]
Device_ID	The identifier of the host system device [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data read and written per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
reads/s	The total number of read requests per second to the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Storage_Path	The name of the storage path [Sequential = ID Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
throughputContention	The average amount of time in milliseconds for an I/O operation to complete. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = AVG]
throughputUsage	The average amount of data transferred per second in kilobytes by the storage path. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
writes/s	The total number of write requests per second to the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]

17.1.43. VMware Storage.Path Summary Table

The VMware Storage.Path Summary table stores I/O operation data summarized by storage paths.

Table Field Hierarchy

Class:	VMware Storage
Subclass:	Path Summary
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	VMware Storage.Path Summary
Open Table Name:	VMSTORPATHSUMMARY
Collection interval:	Based on the collection period
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	Storage Path Summary

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
highestLatency	The highest latency in milliseconds of the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

17.1.44. VMware Storage.Virtual Disk by Virtual Machine Table

The VMware Storage.Virtual Disk by Virtual Machine table stores I/O operation data on virtual disks by virtual machine.

Table Field Hierarchy

Class:	VMware Storage
Subclass:	Virtual Disk by Virtual Machine
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
TeamQuest Table Name:	VMware Storage.Virtual Disk by Virtual Machine
Open Table Name:	VMSTORVIRTDISKBYVM
Collection interval:	Based on the collection period
Default retention:	3 days at collection period interval 8 days at 10-minute intervals 4 months at 1-hour intervals 9 months at 24-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of the virtual disk. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of the virtual disk. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
busRst/s	The number of SCSI bus resets per second that occurred on the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
capacity	The capacity of the virtual disk in gigabytes. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = LST Non-Sequential = SUM]
cmds/s	The number of commands (requests) issued per second to the virtual disk. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]

cmdsAbrt/s	The number of SCSI commands aborted by the storage adapter per second. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data read and written per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
readLoadMetric	The storage DRS virtual disk statistic for the read workload model. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
readOIO	The number of outstanding read requests to the virtual disk during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
reads/s	The total number of read requests per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
throughputContention	The average amount of time in milliseconds taken by an I/O operation to complete. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = AVG]
throughputUsage	The average amount of data transferred per second in kilobytes by the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
Virtual_Disk	The name of the virtual disk [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

writeLoadMetric	The storage DRS virtual disk statistic for the write workload model. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
writeOIO	The number of outstanding write requests to the virtual disk during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
writes/s	The total number of write requests per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]

17.2. Derived Tables

A derived table definition allows multiple stored tables to be brought together into a single logical table. A derived table is created by indicating which fields from one or more stored tables are brought together, and which fields are to be calculated based on values of other fields in the same record when the stored tables have been brought together.

A derived table definition can reference one or more stored tables. A derived table cannot reference another derived table. One of the stored tables referenced by a derived table must be identified as the *primary reference table*. The primary reference table is the table that the other tables are joined to. The other reference tables are referred to as the *secondary reference tables*. One of the reference tables can be flagged as the table to use when determining the selection statements for applying an IT Resource to the derived table. This table is referred to as the *primary selection table*.

The following tables are derived tables and therefore are only available for viewing in TeamQuest Analyzer:

In this subsection, you can find a listing of the derived table statistics collected by the agent:

- Cluster CPU Summary Table (see 17.2.1)
- Cluster Memory Summary Table (see 17.2.2)
- Cluster Resource Allocation Table (see 17.2.3)
- Cluster Virtual Machine Operations Table (see 17.2.4)
- Datastore Cluster File Type Usage Table (see 17.2.5)
- Datastore Cluster Summary Table (see 17.2.6)
- Datastore Cluster Usage by Virtual Machine Table (see 17.2.7)
- Host Block Device Summary Table (see 17.2.8)
- Host Block Device Usage Table (see 17.2.9)
- Host CPU Resource Usage Table (see 17.2.10)
- Host CPU Summary Table (see 17.2.11)
- Host CPU Usage Table (see 17.2.12)

- Host Memory Summary Table (see 17.2.13)
- Host Network Device Summary Table (see 17.2.14)
- Host Network Device Usage Table (see 17.2.15)
- Resource Pool CPU Summary Table (see 17.2.16)
- Resource Pool Memory Summary Table (see 17.2.17)
- Resource Pool Resource Allocation Table (see 17.2.18)
- Storage Adapter by Host System Table (see 17.2.19)
- Storage Adapter Summary Table (see 17.2.20)
- Storage Datastore by Host System Table (see 17.2.21)
- Storage Datastore by Virtual Machine Table (see 17.2.22)
- Storage Datastore Summary Table (see 17.2.23)
- Storage Path by Host System Table (see 17.2.24)
- Storage Path Summary Table (see 17.2.25)
- Virtual Disk by Virtual Machine Table (see 17.2.26)
- Virtual Machine Availability Table (see 17.2.27)
- Virtual Machine Block Device Usage Table (see 17.2.28)
- Virtual Machine Configuration Table (see 17.2.29)
- Virtual Machine CPU Usage Table (see 17.2.30)
- Virtual Machine Memory Usage Table (see 17.2.31)
- Virtual Machine Network Device Usage Table (see 17.2.32)
- Virtual Machine Virtual CPU Usage Table (see 17.2.33)
- Virtual Machine Network vmnic Usage Table (see 17.2.34)
- VMware CPU Relative Performance Table (see 17.2.35)

17.2.1. Cluster CPU Summary Table

The Cluster CPU Summary table is derived from the VMware Cluster.CPU Summary and the VMware.Host Configuration tables. You can view the cluster CPU summary data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Cluster CPU Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Cluster/ <i>Cluster(ClusterId)</i>
Table type:	Performance
Physical tables used to produce this table:	VMware Cluster.CPU Summary VMware.Host Configuration

Statistic Name	Description
%busy	The percentage of elapsed CPU time the processors were busy across all of the virtual machines in the VMware cluster. A value of <N/A> is displayed when the VMware Distributed Resource Scheduler (DRS) is disabled for the VMware cluster. [Sequential = AVG Non-Sequential = AVG]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the VMware cluster to which the data applies [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a VMware cluster is renamed. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]

effectivecpu	<p>The total amount of available CPU in megahertz (MHz) of all of the hosts within a VMware cluster. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. Calculated as</p> $\text{effectivecpu} = \text{aggregate host CPU capacity} - (\text{VMKernel} + \text{service console CPU} + \text{other service CPU})$ <p>[Sequential = AVG Non-Sequential = AVG]</p>
Interval	<p>The expected sampling interval in seconds</p> <p>[Sequential = SUM Non-Sequential = AVG]</p>
reservedCapacity	<p>The total CPU capacity in megahertz (MHz) reserved by all of the virtual machines within a VMware cluster. This statistic is only available for VMware vCenter server statistic level 2.</p> <p>[Sequential = LST Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
total(MHz)	<p>The total amount of CPU in megahertz (MHz) of all of the hosts within the VMware cluster. The maximum value is equal to the frequency of the processors multiplied by the number of cores. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
usage(MHz)	<p>The CPU usage in megahertz (MHz) of all of the powered-on virtual machines in a VMware cluster. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster.</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
vCenter	<p>The name of the VMware vCenter server</p> <p>[Sequential = ID Non-Sequential = ID]</p>

17.2.2. Cluster Memory Summary Table

The Cluster Memory Summary table is derived from the VMware Cluster.Memory Summary and the VMware.Host Configuration tables. You can view the cluster memory summary data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Cluster Memory Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Cluster/ <i>Cluster(ClusterId)</i>
Table type:	Performance
Physical tables used to produce this table:	VMware Cluster.Memory Summary VMware.Host Configuration

Statistic Name	Description
%usage	The percentage of total available memory that is used. A value of <N/A> is displayed when the VMware Distributed Resource Scheduler (DRS) is disabled for the VMware cluster. Calculated as $\text{memory usage} = (\text{memory consumed} + \text{memory overhead}) / \text{effectivemem}$ [Sequential = AVG Non-Sequential = AVG]
active	The working set size estimate in megabytes at the end of the interval for the VMware cluster. This statistic is only available for VMware vCenter server statistic level 2. [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
balloon	The amount of memory in megabytes allocated by the virtual machine memory control driver. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. [Sequential = AVG Non-Sequential = SUM]
consumed	The amount of host memory in megabytes consumed by all of the powered-on virtual machines for guest memory within a VMware cluster [Sequential = AVG Non-Sequential = SUM]
Cluster	The name of the VMware cluster to which the data applies [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a VMware cluster is renamed. [Sequential = ID Non-Sequential = ID]

Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
effectivemem	The total amount of memory in megabytes of all of the hosts within a VMware cluster that are available for the virtual machine memory and the virtual machine overhead memory. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. [Sequential = AVG Non-Sequential = AVG]
granted	The amount of all granted memory in megabytes for all of the powered-on virtual machines. This statistic is only available for VMware vCenter server statistic level 2. [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
overhead	The amount of additional host memory in megabytes allocated to the virtual machine. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. [Sequential = AVG Non-Sequential = SUM]
reservedCapacity	The amount of memory in megabytes reserved by the virtual machines on a VMware host. This statistic is only available for VMware vCenter server statistic level 2. [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shared	The amount of guest physical memory in megabytes shared with other virtual machines. This value includes the amount of zero memory. [Sequential = AVG Non-Sequential = SUM]
swapused	The amount of memory in megabytes that is used by swap. This statistic is only available for VMware vCenter server statistic level 2. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. [Sequential = AVG Non-Sequential = SUM]
sysUsage	The amount of memory in megabytes used by the VMKernel for core functionality, such as device drivers and other internal usage components. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. This statistic is only available for VMware vCenter server statistic level 2 and for VMware vCenter Server 4.0 and lower. [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]

total	The total amount of memory in megabytes of all of the hosts within a VMware cluster that are available for the virtual machine memory and virtual machine overhead memory. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
zero	The amount of memory in megabytes that only contains 0 values. It is included in shared memory. Through transparent page sharing, zero memory pages can be shared among virtual machines that run the same operating system. A value of <N/A> is displayed when the VMware DRS is disabled for the VMware cluster. [Sequential = AVG Non-Sequential = SUM]

17.2.3. Cluster Resource Allocation Table

The Cluster Resource Allocation table is derived from the VMware Cluster.Resource Allocation and the VMware.Host Configuration tables. You can view the cluster resource allocation data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Cluster Resource Allocation
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/cluster/Cluster(ClusterId)
Table type:	Performance
Physical tables used to produce this table:	VMware Cluster.Resource Allocation VMware.Host Configuration

Statistic Name	Description
activeHosts	The number of active hosts in the VMware cluster [Sequential = LST Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the data applies [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. [Sequential = ID Non-Sequential = ID]

cpuAvailable	The total amount of CPU in megahertz (MHz) available to satisfy a reservation for all of the virtual machines and resource pools in the cluster [Sequential = LST Non-Sequential = SUM]
cpuReserved	The total amount of CPU in megahertz (MHz) that has been used to satisfy the reservation requirements of all of the descendants of the virtual machines and resource pools in the cluster [Sequential = LST Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
memAvailable	The total amount of memory in megabytes that is available to satisfy the reservation for all of the virtual machines and resource pools in the cluster [Sequential = LST Non-Sequential = SUM]
memOverhead	The total amount of memory in megabytes that has been used to satisfy the reservation requirements of all of the descendants of the running virtual machines in the cluster [Sequential = LST Non-Sequential = SUM]
memReserved	The total amount of memory in megabytes that has been used to satisfy the reservation requirements of all of the descendants of the virtual machines and resource pools in the cluster [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totalCpu	The total available CPU in megahertz (MHz) of all of the hosts within the cluster [Sequential = LST Non-Sequential = SUM]
totalMem	The total amount of memory in megabytes of all of the hosts within the cluster that is available for use for virtual machine memory and virtual machine overhead memory [Sequential = LST Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.4. Cluster Virtual Machine Operations Table

The Cluster Virtual Machine Operations table is derived from the VMware Cluster.Virtual Machine Operations and the VMware.Host Configuration tables. You can view the cluster virtual machine operations data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Cluster Virtual Machine Operations
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Cluster/ <i>Cluster(ClusterId)</i>
Table type:	Performance
Physical tables used to produce this table:	VMware Cluster.Virtual Machine Operations VMware.Host Configuration

Statistic Name	Description
activeVMs	The number of active virtual machines in the VMware cluster [Sequential = LST Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
clone	The number of virtual machine clone operations [Sequential = SUM Non-Sequential = SUM]
Cluster	The name of the cluster to which the data applies [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. [Sequential = ID Non-Sequential = ID]
create	The number of virtual machine create operations [Sequential = SUM Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
datastoreChange	The number of datastore change operations for powered-off and suspended virtual machines [Sequential = SUM Non-Sequential = SUM]
delete	The number of virtual machine delete operations [Sequential = SUM Non-Sequential = SUM]
failover	The number of virtual machine failover operations. This statistic is only available for clusters with VMware High Availability (HA) on. [Sequential = LST Non-Sequential = SUM]

guestReboot	The number of virtual machine guest reboot operations [Sequential = SUM Non-Sequential = SUM]
guestShutdown	The number of virtual machine guest shutdown operations [Sequential = SUM Non-Sequential = SUM]
hostChange	The total number of host change operations for powered-down and suspended virtual machines [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
powerOff	The number of virtual machine power-off operations [Sequential = SUM Non-Sequential = SUM]
powerOn	The number of virtual machine power-on operations [Sequential = SUM Non-Sequential = SUM]
reconfigure	The number of virtual machine reconfigure operations [Sequential = SUM Non-Sequential = SUM]
register	The number of virtual machine register operations [Sequential = SUM Non-Sequential = SUM]
reset	The number of virtual machine reset operations [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
standByGuest	The number of virtual machine standby guest operations [Sequential = SUM Non-Sequential = SUM]
storageMotion	The number of migrations with storage vMotion or datastore change operations for all of the powered-on virtual machines [Sequential = SUM Non-Sequential = SUM]
suspend	The number of virtual machine suspend operations [Sequential = SUM Non-Sequential = SUM]
templateDeploy	The number of virtual machine template deploy operations [Sequential = SUM Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
unregister	The number of virtual machine unregister operations [Sequential = SUM Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
vMotion	The number of migrations with vMotion for powered-on virtual machines [Sequential = SUM Non-Sequential = SUM]

17.2.5. Datastore Cluster File Type Usage Table

The Datastore Cluster File Type Usage table is derived from the VMware.Datastore File Type Usage and the VMware.Storage Configuration tables. You can view the datastore summary data by datacenter, by datastore, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Datastore Cluster File Type Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Datastore/Datastore
Table type:	Performance
Physical tables used to produce this table:	VMware.Datastore.File Type Usage by Datacenter VMware.Storage Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
capacity	The configured size of the VMware datastore in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]
Datastore_Cluster	The name of the VMware datastore cluster to which the data applies [Sequential = ID Non-Sequential = ID]
Datastore_ClusterId	The identifier assigned to the VMware datastore cluster that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
free	The amount of free physical VMware datastore space in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
other	The amount of disk space used by all other non-virtual machine files in gigabytes (GB), such as documentation files and backup files [Sequential = LST Non-Sequential = SUM]

other_vm_files	The amount of disk space used by all other virtual machine files [Sequential = ID Non-Sequential = ID]
provisioned	The amount of physical space in gigabytes (GB) provisioned by an administrator for the VMware datastore. This value is the maximum storage size to which files on the VMware datastore can grow. [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
snapshots	The amount of disk space in gigabytes (GB) used by virtual machine snapshot files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. The virtual machine snapshot files have the extension of .vmsn. [Sequential = LST Non-Sequential = SUM]
swapfiles	The amount of disk space in gigabytes (GB) used by swap files. Swap files are used to back up the virtual machine physical memory. [Sequential = LST Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
used	The amount of physical VMware datastore space in use in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
virtual_disks	The amount of disk space in gigabytes (GB) used by virtual disk files and delta disk files. Virtual disk files store the contents of the virtual machines hard disk drive, including information that is written to virtual machines hard disk (the operating system, program files, and data files). The virtual disk files have the file extension of .vmdk and appear as a physical disk drive on a guest operating system. Delta disk files store the updates made by the virtual machine to the virtual disks after a snapshot is taken. [Sequential = LST Non-Sequential = SUM]

17.2.6. Datastore Cluster Summary Table

The Datastore Cluster Summary table is derived from the VMware.Datastore Summary and the VMware.Storage Configuration tables. You can view the datastore summary data by datacenter, by datastore, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Datastore Cluster Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Datastore/Datastore
Table type:	Performance
Physical tables used to produce this table:	VMware.Datastore.Summary VMware.Storage.Configuration

Statistic Name	Description
%free	The percentage of free physical VMware datastore space [Sequential = LST Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
capacity	The configured size of the VMware datastore in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]
Datastore_Cluster	The name of the VMware datastore cluster to which the data applies [Sequential = ID Non-Sequential = ID]
Datastore_ClusterId	The identifier assigned to the VMware datastore cluster that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
free	The amount of free physical VMware datastore space in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
used	The amount of physical VMware datastore space in use in gigabytes (GB) [Sequential = LST Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.7. Datastore Cluster Usage by Virtual Machine Table

The Datastore Cluster Usage by Virtual Machine table is derived from the VMware.Datastore Usage by Virtual Machine and the VMware.Storage Configuration tables. You can view the datastore summary data by datacenter, by datastore, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Datastore Cluster Usage by Virtual Machine
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Datastore/Datastore
Table type:	Performance
Physical tables used to produce this table:	VMware.Datastore.Usage by Virtual Machine VMware.Storage Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]
Datastore_Cluster	The name of the VMware datastore cluster to which the data applies [Sequential = ID Non-Sequential = ID]

Datastore_ClusterId	The identifier assigned to the VMware datastore cluster that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
Host	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
provisioned	The amount of storage space in gigabytes (GB) provisioned by an administrator for the virtual machine [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
unshared	The amount of storage space in gigabytes (GB) associated exclusively with a virtual machine [Sequential = LST Non-Sequential = SUM]
used	The amount of physical VMware datastore space in use in gigabytes (GB) by the virtual machine [Sequential = LST Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

17.2.8. Host Block Device Summary Table

The Host Block Device Summary table is derived from the Block Device.VMware Summary and the VMware.Host Configuration tables. You can view the host block device summary data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Host Block Device Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	Block Device.VMware Summary VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
cmds/s	The total number of read and write command requests per second [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]

KB/s	The amount of data read and written per second in kilobytes (KB) for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]
reads/s	The total number of read requests per second for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
writes/s	The total number of write requests per second for all of the disk instances of the host system [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.9. Host Block Device Usage Table

The Host Block Device Usage table is derived from the Block Device.by Host System Device and the VMware.Host Configuration tables. You can view the host block device usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Host Block Device Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	Block Device.by Host System Device VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgDeviceLatency	The average amount of time in milliseconds taken to complete a command to the physical device [Sequential = AVG Non-Sequential = AVG]
avgDeviceReadLatency	The average amount of time in milliseconds taken to complete a read operation from the physical device [Sequential = AVG Non-Sequential = AVG]
avgDeviceWriteLatency	The average amount of time in milliseconds taken to complete a write to the physical device [Sequential = AVG Non-Sequential = AVG]
avgKernelLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per command [Sequential = AVG Non-Sequential = AVG]
avgKernelReadLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per read [Sequential = AVG Non-Sequential = AVG]
avgKernelWriteLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per write [Sequential = AVG Non-Sequential = AVG]
avgLatency	The average amount of time in milliseconds taken to complete a command request (queue and disk service time) by the host system disk [Sequential = AVG Non-Sequential = AVG]

avgQueueLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per command [Sequential = AVG Non-Sequential = AVG]
avgQueueReadLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per read [Sequential = AVG Non-Sequential = AVG]
avgQueueWriteLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per write [Sequential = AVG Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of a guest operating system [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of a guest operating system [Sequential = AVG Non-Sequential = AVG]
busRst/s	The number of bus resets per second that occurred on the host system disk [Sequential = AVG Non-Sequential = SUM]
cmds/s	The number of commands (requests) issued per second to the host system disk [Sequential = AVG Non-Sequential = SUM]
cmdsAbrt/s	The number of commands aborted by the host system disk per second [Sequential = AVG Non-Sequential = SUM]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the host system disk [Sequential = AVG Non-Sequential = SUM]

KB_write/s	The amount of data written per second in kilobytes (KB) by the host system disk [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data transferred per second in kilobytes (KB) by the host system disk [Sequential = AVG Non-Sequential = SUM]
reads/s	The number of read requests issued per second to the host system disk [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the disk device [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shares	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = LST Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totalTime	The total time in milliseconds for all command requests on a VMware host [Sequential = SUM Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
writes/s	The number of write requests issued per second to the host system disk [Sequential = AVG Non-Sequential = SUM]

17.2.10. Host CPU Resource Usage Table

The Host CPU Resource Usage table is derived from the CPU.by VMware Resource and the VMware.Host Configuration tables. You can view the host CPU resource usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Host CPU Resource Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	CPU.by VMware Resource VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Object	The name of the object for the host system [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]

System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
usage(MHz)	The CPU usage in megahertz (MHz) over the collection interval [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.11. Host CPU Summary Table

The Host CPU Summary table is derived from the CPU.VMware Summary and the VMware.Host Configuration tables. You can view the host CPU summary data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Host CPU Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	CPU.VMware Summary VMware.Host Configuration

Statistic Name	Description
%busy	The percentage of the CPU used [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]

DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
idle	The amount of processor time in seconds that is spent in an idle state [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
online_cpus	The number of logical CPUs that were online [Sequential = LST Non-Sequential = SUM]
online_cpus_physical	The number of physical CPUs that were online [Sequential = LST Non-Sequential = SUM]
reservedCapacity	The total CPU capacity in megahertz (MHz) reserved by all of the virtual machines [Sequential = LST Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
uptime_t	The total time in days elapsed since the last VMware host reboot [Sequential = LST Non-Sequential = SUM]
usage(MHz)	The CPU usage in megahertz (MHz) over the collected interval [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.12. Host CPU Usage Table

The Host CPU Usage table is derived from the CPU.by Host Processor and the VMware.Host Configuration tables. You can view the host CPU usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Host CPU Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	CPU.by Host Processor VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the data applies [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
idle	The amount of processor time in seconds that is spent in an idle state [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Object	The name of the CPU object [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
usage	The percentage of time the CPU is in use over the collection interval [Sequential = AVG Non-Sequential = SUM]
usedsec	The processor time in seconds consumed by the VMware host [Sequential = SUM Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.13. Host Memory Summary Table

The Host Memory Summary table is derived from the Memory.VMware Summary and the VMware.Host Configuration tables. You can view the host memory summary data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Host Memory Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	Memory.VMware Summary VMware.Host Configuration

Statistic Name	Description
%usage	The percentage of memory usage over the collection interval [Sequential = AVG Non-Sequential = SUM]
active	The working set size estimate in megabytes at the end of the interval for the host [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]

consumed	The amount of memory in megabytes used on the host. It includes service console memory, VMKernel memory, VMware Infrastructure services memory, and VM memory. [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
free_mem	The amount of memory in megabytes currently available to be used by the host system [Sequential = LST Non-Sequential = SUM]
heap	The amount of memory in megabytes allocated for the heap [Sequential = AVG Non-Sequential = SUM]
heapfree	The amount of free space in megabytes in the memory heap [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
managed_mem	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = LST Non-Sequential = SUM]
overhead	The amount of additional host memory in megabytes allocated to the virtual machine [Sequential = AVG Non-Sequential = SUM]
reservedCapacity	The amount of memory in megabytes reserved by the virtual machines on a VMware host [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shared	The amount of memory in megabytes shared between the virtual machines on a VMware host [Sequential = AVG Non-Sequential = ID]
shared_comm_mem	The total amount of shared common memory in megabytes on the host system [Sequential = LST Non-Sequential = SUM]
size	The amount of memory in megabytes granted [Sequential = AVG Non-Sequential = SUM]
state	The VMKernel threshold for the amount of free memory on the host [Sequential = LST Non-Sequential = SUM]
swapin	The total amount of memory in megabytes that is swapped in on a VMware host [Sequential = AVG Non-Sequential = SUM]
swapout	The total amount of memory in megabytes that is swapped out on a VMware host [Sequential = AVG Non-Sequential = SUM]

swapused	The amount of memory in megabytes that is used by swap [Sequential = AVG Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
sysUsage	The amount of memory in megabytes used by the VMKernel for core functionality, such as device drivers and other internal usage [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
total_mem	The total amount of physical memory in megabytes on the host system [Sequential = LST Non-Sequential = SUM]
unreserved	The amount of memory in megabytes that is unreserved [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
vmmemctl	The amount of memory in megabytes used by memory control [Sequential = LST Non-Sequential = SUM]
zero	The amount of memory in megabytes that is zeroed out [Sequential = AVG Non-Sequential = SUM]

17.2.14. Host Network Device Summary Table

The Host Network Device Summary table is derived from the Network Device.VMware Summary and the VMware.Host Configuration tables. You can view the host network device summary data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Host Network Device Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	Network Device.VMware Summary VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KBRx/s	The amount of data received per second in kilobytes (KB) [Sequential = AVG Non-Sequential = SUM]
KBTx/s	The amount of data transmitted per second in kilobytes (KB) [Sequential = AVG Non-Sequential = SUM]
KBx/s	The amount of data transferred per second in kilobytes (KB) [Sequential = AVG Non-Sequential = SUM]
packets/s	The total number of packets transferred per second [Sequential = AVG Non-Sequential = SUM]

pktsRx/s	The total number of packets received per second [Sequential = AVG Non-Sequential = SUM]
pktsTx/s	The total number of packets transmitted per second [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.15. Host Network Device Usage Table

The Host Network Device Usage table is derived from the Network Device.by Host System Device and the VMware.Host Configuration tables. You can view the host network device usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Host Network Device Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	Network Device.by Host System Device VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KBRx/s	The amount of data received per second in kilobytes (KB) by the host system interface [Sequential = AVG Non-Sequential = SUM]
KBTx/s	The amount of data transmitted per second in kilobytes (KB) by the host system interface [Sequential = AVG Non-Sequential = SUM]
KBx/s	The amount of data transferred per second in kilobytes (KB) by the host system interface [Sequential = AVG Non-Sequential = SUM]

packets/s	The total number of packets transferred per second by the host system interface [Sequential = AVG Non-Sequential = SUM]
pktsRx/s	The number of packets received per second by the host system interface [Sequential = AVG Non-Sequential = SUM]
pktsTx/s	The number of packets transmitted per second by the host system interface [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the host system interface [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.16. Resource Pool CPU Summary Table

The Resource Pool CPU Summary table is derived from the VMware Resource Pool.CPU Summary and the VMware.Host Configuration tables. You can view the resource pool CPU summary data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Resource Pool CPU Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/ResourcePool/Resource_pool(Resource Pool_Id)
Table type:	Performance
Physical tables used to produce this table:	VMware Resource Pool.CPU Summary VMware.Host Configuration

Statistic Name	Description
%busy	The percentage of elapsed CPU time the processors were busy across all of the virtual machines in the resource pool [Sequential = AVG Non-Sequential = AVG]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
capacityContention	The percentage of time the virtual machine is unable to run because it is contending for access to the physical CPUs. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
capacityDemand	The amount of CPU resources a virtual machine would use if there were no CPU contention or CPU limit. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
capacityEntitlement	The amount of CPU resources devoted by the VMware ESX scheduler to virtual machines and resource pools. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
Cluster	The name of the cluster to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware host. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]

corecountContention	The amount of time the virtual machine was ready to run but was unable to run due to co-scheduling constraints. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
corecountProvisioned	The number of virtual processors or physical cores provisioned to the entity. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Host	The name of the host system to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware cluster. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]
ResourcePool	The name of the resource pool to which the data applies [Sequential = ID Non-Sequential = ID]
ResourcePool_Id	The unique identifier for a resource pool provided by the VMware API. This identifier does not change when a resource pool is renamed. [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
usage(MHz)	The CPU usage in megahertz (MHz) over the collection interval [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.17. Resource Pool Memory Summary Table

The Resource Pool Memory Summary table is derived from the VMware Resource Pool.Memory Summary and the VMware.Host Configuration tables. You can view the resource pool memory summary data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Resource Pool Memory Summary
Subclass:	
IT Resource Name:	<i>/TeamQuest/System/VMware/ResourcePool/Resource_pool(Resource Pool_Id)</i>
Table type:	Performance
Physical tables used to produce this table:	VMware Resource Pool.Memory Summary VMware.Host Configuration

Statistic Name	Description
%usage	The percentage of memory usage over the collection interval. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = AVG]
active	The working set size estimate in megabytes at the end of the interval for the resource pool [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
balloon	The amount of memory in megabytes allocated by the virtual machine memory control driver (vmmemctl), which is installed with VMware Tools. The vmmemctl is a memory management driver that controls ballooning. [Sequential = AVG Non-Sequential = SUM]
capacityContention	The percentage of time the virtual machine is waiting to access swapped or compressed memory. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
capacityEntitlement	The amount of host physical memory devoted by the VMware ESX scheduler to the virtual machine. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
capacityUsage	The amount of physical memory in megabytes used by the virtual machine. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]

Cluster	The name of the cluster to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware host. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
compressed	The amount of memory in megabytes compressed by the VMware ESX Server. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
compressionRate	The rate of memory compression for the virtual machine. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
consumed	The amount of memory in megabytes consumed by a resource pool, host, or virtual machine [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
decompressionRate	The rate of memory decompression for the virtual machine. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]
granted	The amount of all granted memory in megabytes for all of the powered-on virtual machines. This statistic is only available for VMware vCenter server statistic level 2. [Sequential = AVG Non-Sequential = SUM]
Host	The name of the host system to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware cluster. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]
overhead	The amount of additional resource pool memory in megabytes allocated to the virtual machine [Sequential = AVG Non-Sequential = SUM]
ResourcePool	The name of the resource pool to which the data applies [Sequential = ID Non-Sequential = ID]
ResourcePool_Id	The unique identifier for a resource pool provided by the VMware API. This identifier does not change when a resource pool is renamed. [Sequential = ID Non-Sequential = ID]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shared	The amount of virtual machine memory in megabytes that is shared with other virtual machines, relative to a single virtual machine or to all powered-on virtual machines on a host [Sequential = AVG Non-Sequential = SUM]
swapped	The amount of memory in megabytes currently swapped to the VMware File System 3 (VMFS3) swap file [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
zero	The amount of memory in megabytes that is zeroed out. This statistic is only available for VMware vCenter Server 5.0 and higher. [Sequential = AVG Non-Sequential = SUM]

17.2.18. Resource Pool Resource Allocation Table

The Resource Pool Resource Allocation table is derived from the VMware Resource Pool.Resource Allocation and the VMware.Host Configuration tables. You can view the resource pool resource allocation data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Resource Pool Resource Allocation
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/ResourcePool/ <i>Resource_Pool_Id</i>
Table type:	Performance
Physical tables used to produce this table:	VMware Resource Pool.Resource Allocation VMware.Host Configuration

Statistic Name	Description
activeVMs	The number of active virtual machines in a resource pool [Sequential = LST Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
childResourcePool	The number of child resource pools in a resource pool [Sequential = LST Non-Sequential = SUM]
Cluster	The name of the cluster to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware host. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
cpuAvailable	The total amount of CPU in megahertz (MHz) that is available to satisfy a reservation requirement for all virtual machines in a resource pool [Sequential = LST Non-Sequential = LST]
cpuReserved	The total amount of CPU in megahertz (MHz) that is used to satisfy the reservation requirements of all descendants of the resource pools and virtual machines in a resource pool [Sequential = LST Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]

DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Host	The name of the host system to which the resource pool belongs. This field contains <N/A> if the root resource pool is a VMware cluster. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Sequential = ID Non-Sequential = ID]
memAvailable	The total amount of memory in megabytes that is available to satisfy the reservation requirements for all virtual machines in a resource pool [Sequential = LST Non-Sequential = LST]
memOverhead	The total amount of memory in megabytes used to satisfy the reservation requirements of all descendants of the virtual machines in a resource pool or any of the child resource pools [Sequential = LST Non-Sequential = SUM]
memReserved	The total amount of memory in megabytes used to satisfy the reservation requirements of all descendants of the resource pools and virtual machines in a resource pool [Sequential = LST Non-Sequential = SUM]
ResourcePool	The name of the resource pool to which the data applies [Sequential = ID Non-Sequential = ID]
ResourcePool_Id	The unique identifier for a resource pool provided by the VMware API. This identifier does not change when a resource pool is renamed. [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.19. Storage Adapter by Host System Table

The Storage Adapter by Host System table is derived from the VMware Storage.Adapter by Host System and the VMware.Host Configuration tables. You can view the storage adapter availability data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Storage Adapter by Host System
Subclass:	
IT Resource Name:	/TeamQuest/System/
Table type:	Performance
Physical tables used to produce this table:	VMware.Storage.Adapter by Host System VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgQueueLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per command from the perspective of a storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of a storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of a storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
cmds/s	The number of commands (requests) issued per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]

DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data transferred per second in kilobytes (KB) by the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
oIOsPct	The percentage of I/O operations that have been issued but have not yet completed. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
outstandingIOs	The number of I/O operations that have been issued but have not yet completed. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
queued	The number of I/O operations waiting to be issued. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
queueDepth	The maximum number of I/O operations that can be outstanding at a time. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
reads/s	The number of read requests issued per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Storage_Adapter	The name of the storage adapter [Sequential = ID Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
throughputContention	The average amount of time in milliseconds for an I/O operation to complete. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = AVG]
throughputUsage	The average amount of data transferred per second in kilobytes by the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
writes/s	The total number of write requests issued per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.20. Storage Adapter Summary Table

The Storage Adapter Summary table is derived from the VMware Storage.Adapter Summary and the VMware.Host Configuration tables. You can view the storage adapter availability data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Storage Adapter Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/
Table type:	Performance
Physical tables used to produce this table:	VMware Storage.Adapter Summary VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
cmds/s	The number of commands (requests) issued per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]

highestLatency	The highest latency in milliseconds of the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = MAX Non-Sequential = MAX]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.21. Storage Datastore by Host System Table

The Storage Datastore by Host System table is derived from the VMware Storage.Datastore by Host System and the VMware.Storage Configuration tables.

Table Field Hierarchy

Class:	Storage Datastore by Host System
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	VMware Storage.Datastore by Host System VMware.Storage Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgLatency	The average amount of time in milliseconds taken to complete a command request (queue and disk service time) by the host system disk [Sequential = AVG Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the datastore [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the datastore [Sequential = AVG Non-Sequential = AVG]

cmds	The number of commands (requests) issued to the datastore during the collection interval. Calculated as cmds= reads + writes [Sequential = SUM Non-Sequential = SUM]
cmds/s	The number of commands (requests) issued per second to the datastore. Calculated as cmds/s= reads/s + writes/s [Sequential = SUM Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the datastore belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
datastoreIops	The aggregated number of storage I/O operations on the datastore [Sequential = AVG Non-Sequential = SUM]
datastoreMaxQueueDepth	The maximum number of storage I/O operations supported by the datastore [Sequential = LST Non-Sequential = LST]
datastoreNormalReadLatency	The storage DRS normalized read latency for the datastore [Sequential = LST Non-Sequential = LST]
datastoreNormalWriteLatency	The storage DRS normalized write latency for the datastore [Sequential = LST Non-Sequential = LST]
datastoreReadBytes	The number of storage DRS bytes read by the datastore [Sequential = LST Non-Sequential = LST]
datastoreReadIops	The number of storage DRS read I/O operations of the datastore [Sequential = LST Non-Sequential = LST]
datastoreReadLoadMetric	The storage DRS read workload metric of the datastore [Sequential = LST Non-Sequential = LST]
datastoreReadOIO	The number of outstanding read requests by the datastore [Sequential = LST Non-Sequential = LST]
datastoreWriteBytes	The number of storage DRS bytes written to the datastore [Sequential = LST Non-Sequential = LST]
datastoreWriteIops	The number of storage DRS write I/O operations of the datastore [Sequential = LST Non-Sequential = LST]
datastoreWriteLoadMetric	The storage DRS write workload metric of the datastore [Sequential = LST Non-Sequential = LST]
datastoreWriteOIO	The number of outstanding write requests by the datastore [Sequential = LST Non-Sequential = LST]
highestLatency	The highest latency value across all datastores used by the host [Sequential = MAX Non-Sequential = MAX]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]

KB_read/s	The amount of data read per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data read and written per second in kilobytes (KB) by the datastore. Calculated as $KB/s = KB_reads/s + KB_writes/s$ [Sequential = AVG Non-Sequential = SUM]
reads	The number of read requests issued to the datastore during the collection interval [Sequential = SUM Non-Sequential = SUM]
reads/s	The number of read requests issued per second to the datastore [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
sizeNormalized DatastoreLatency	The normalized latency in microseconds on the datastore. Data for all virtual machines is combined into this statistic. [Sequential = AVG Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
writes	The number of write requests issued to the datastore during the collection interval [Sequential = SUM Non-Sequential = SUM]
writes/s	The number of write requests issued per second to the datastore [Sequential = AVG Non-Sequential = SUM]

17.2.22. Storage Datastore by Virtual Machine Table

The Storage Datastore by Virtual Machine table is derived from the VMware Storage.Datastore by Virtual Machine and the VMware.Storage Configuration tables.

Table Field Hierarchy

Class:	Storage Datastore by Virtual Machine
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Performance
Physical tables used to produce this table:	VMware.Storage.Datastore by Virtual Machine VMware.Storage Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the datastore [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the datastore [Sequential = AVG Non-Sequential = AVG]
cmds	The number of commands (requests) issued per second to the datastore. Calculated as cmds= reads + writes [Sequential = SUM Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the datastore belongs [Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Non-Sequential = ID]
highestLatency	The highest latency value across all datastores used by the host [Sequential = MAX Non-Sequential = MAX]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]

KB/s	<p>The amount of data requests per second in kilobytes (KB) by the datastore. Calculated as</p> $\text{KB/s} = \text{KB_reads} + \text{KB_writes}$ <p>[Sequential = AVG Non-Sequential = SUM]</p>
reads	<p>The number of read requests issued to the datastore during the collection interval</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
reads/s	<p>The number of read requests issued per second to the datastore</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>
Sample_End_Time	<p>The timestamp of the actual end of data collection for the current sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
System	<p>The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated.</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Time	<p>The timestamp of the data sample</p> <p>[Sequential = LST Non-Sequential = ID]</p>
UniqueId	<p>The identifier assigned to a datastore that is unique within a VMware vCenter server</p> <p>[Sequential = ID Non-Sequential = ID]</p>
Virtual_Machine	<p>The name of the virtual machine to which the data applies</p> <p>[Sequential = ID Non-Sequential = ID]</p>
writes	<p>The number of write requests issued per second to the datastore</p> <p>[Sequential = SUM Non-Sequential = SUM]</p>
writes/s	<p>The number of write requests issued to the datastore during the collection interval</p> <p>[Sequential = AVG Non-Sequential = SUM]</p>

17.2.23. Storage Datastore Summary Table

The Storage Datastore Summary table is derived from the VMware Storage.Datastore Summary and the VMware.Storage Configuration tables.

Table Field Hierarchy

Class:	Storage Datastore Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Datastore/Datastore
Table type:	Performance
Physical tables used to produce this table:	VMware.Storage.Datastore Summary VMware.Storage Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Datacenter	The name of the datacenter to which the datastore belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Datastore	The name of the VMware datastore to which the data applies [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the datastore [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data read and written per second in kilobytes (KB) by the datastore. Calculated as KB/s= KB_reads/s + KB_writes/s [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
UniqueId	The identifier assigned to a datastore that is unique within a VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.24. Storage Path by Host System Table

The Storage Path by Host System table is derived from the VMware Storage.Path by Host System and the VMware.Host Configuration tables. You can view the storage path availability data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Storage Path by Host System
Subclass:	
IT Resource Name:	/TeamQuest/System/
Table type:	Performance
Physical tables used to produce this table:	VMware Storage.Path by Host System VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
busRst/s	The number of SCSI bus resets per second that occurred on the storage path. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
cmds/s	The number of commands (requests) issued per second to the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
cmdsAbrt/s	The number of SCSI commands aborted by the storage path per second. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]

Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Device_Name	The name of the host system device [Sequential = ID Non-Sequential = ID]
Device_ID	The identifier of the host system device [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data read and written per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
reads/s	The total number of read requests per second to the storage path. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
Storage_Path	The name of the storage path [Sequential = ID Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
throughputContention	The average amount of time in milliseconds for an I/O operation to complete. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = AVG]
throughputUsage	The average amount of data transferred per second in kilobytes by the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
writes/s	The total number of write requests per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.25. Storage Path Summary Table

The Storage Path Summary table is derived from the VMware Storage.Path Summary and the VMware.Host Configuration tables. You can view the storage path availability data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Storage Path Summary
Subclass:	
IT Resource Name:	/TeamQuest/System/
Table type:	Performance
Physical tables used to produce this table:	VMware.Storage.Path Summary VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
highestLatency	The highest latency in milliseconds of the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]

Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

17.2.26. Virtual Disk by Virtual Machine Table

The Virtual Disk by Virtual Machine table is derived from the VMware Storage.Virtual Disk by Virtual Machine and the VMware.Host Configuration tables. You can view the virtual disk availability data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Disk by Virtual Machine
Subclass:	
IT Resource Name:	/TeamQuest/System/
Table type:	Performance
Physical tables used to produce this table:	VMware Storage.Virtual Disk by Virtual Machine VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of the virtual disk. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of the virtual disk. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = AVG]
busRst/s	The number of SCSI bus resets per second that occurred on the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]

cmdsAbrt/s	The number of SCSI commands aborted by the storage adapter per second. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
KB/s	The amount of data read and written per second in kilobytes (KB). This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
readLoadMetric	The storage DRS virtual disk statistic for the read workload model. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
readOIO	The average number of outstanding read requests to the virtual disk during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
reads/s	The total number of read requests per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
throughputContention	The average amount of time in milliseconds taken by an I/O operation to complete. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = AVG]
throughputUsage	The average amount of data transferred per second in kilobytes by the storage adapter. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = AVG Non-Sequential = SUM]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

Virtual_Disk	The name of the virtual disk [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
writeLoadMetric	The storage DRS virtual disk statistic for the write workload model. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
writeOIO	The average number of outstanding write requests to the virtual disk during the collection interval. This statistic is available for VMware ESX 5.0.0 and later. [Sequential = LST Non-Sequential = SUM]
writes/s	The total number of write requests per second to the storage adapter. This statistic is available for VMware ESX 4.1.0 and later. [Sequential = AVG Non-Sequential = SUM]

17.2.27. Virtual Machine Availability Table

The Virtual Machine Availability table is derived from the VMware.Availability by Virtual Machine and the VMware.Host Configuration tables. You can view the virtual machine availability data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Machine Availability
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Performance
Physical tables used to produce this table:	VMware.Availability by Virtual Machine VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]

DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
heartbeat	The number of heartbeats in the collection period. The heartbeat represents the overall health of the guest operating system. [Sequential = AVG Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
uptime_t	The total time in days elapsed since the last virtual machine reboot [Sequential = LST Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

17.2.28. Virtual Machine Block Device Usage Table

The Virtual Machine Block Device Usage table is derived from the Block Device.by Virtual Machine and the VMware.Host Configuration tables. You can view the virtual machine block device usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Machine Block Device Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Performance
Physical tables used to produce this table:	Block Device.by Virtual Machine VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
avgDeviceLatency	The average amount of time in milliseconds taken to complete a command to the physical device [Sequential = AVG Non-Sequential = AVG]
avgDeviceReadLatency	The average amount of time in milliseconds taken to complete a read from the physical device [Sequential = AVG Non-Sequential = AVG]
avgDeviceWriteLatency	The average amount of time in milliseconds taken to complete a write to the physical device [Sequential = AVG Non-Sequential = AVG]
avgKernelLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per command [Sequential = AVG Non-Sequential = AVG]
avgKernelReadLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per read [Sequential = AVG Non-Sequential = AVG]
avgKernelWriteLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel per write [Sequential = AVG Non-Sequential = AVG]
avgLatency	The average amount of time in milliseconds taken to complete a command request (queue and disk service time) by the host system disk [Sequential = AVG Non-Sequential = AVG]
avgQueueLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per command [Sequential = AVG Non-Sequential = AVG]

avgQueueReadLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per read [Sequential = AVG Non-Sequential = AVG]
avgQueueWriteLatency	The average amount of time in milliseconds spent in the ESX Server VMKernel queue per write [Sequential = AVG Non-Sequential = AVG]
avgReadLatency	The average amount of time in milliseconds taken by a read operation to complete from the perspective of a guest operating system [Sequential = AVG Non-Sequential = AVG]
avgWriteLatency	The average amount of time in milliseconds taken by a write operation to complete from the perspective of a guest operating system [Sequential = AVG Non-Sequential = AVG]
avresp	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. The average response time is now reported under the avgLatency statistic name. [Sequential = AVG Non-Sequential = AVG]
busRst/s	The number of bus resets per second that occurred on the virtual machine [Sequential = AVG Non-Sequential = SUM]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
cmds/s	The number of commands (requests) issued per second to the physical device [Sequential = AVG Non-Sequential = SUM]
cmdsAbrt/s	The number of commands per second that were aborted by the virtual machine [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KB_read/s	The amount of data read per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KB_write/s	The amount of data written per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]

KB/s	The amount of data transferred per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
paeCmds/s	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = AVG Non-Sequential = SUM]
paeCopies/s	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = AVG Non-Sequential = SUM]
reads/s	The number of read requests issued per second to the physical device [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the physical device [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shares	The number of shares allocated to the virtual machine [Sequential = LST Non-Sequential = SUM]
splitCmds/s	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = AVG Non-Sequential = SUM]
splitCopies/s	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = AVG Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
totalTime	The total time in milliseconds for all command requests by the virtual machine [Sequential = SUM Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
writes/s	The total number of write requests issued per second to the physical device [Sequential = AVG Non-Sequential = SUM]

17.2.29. Virtual Machine Configuration Table

The Virtual Machine Configuration table is derived from the VMware.Virtual_Machines and the VMware.Host Configuration tables. You can view the virtual machine configuration data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Machine Configuration
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Event
Physical tables used to produce this table:	VMware.Virtual_Machines VMware.Host Configuration

Statistic Name	Description
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Non-Sequential = ID]
Configuration_File	The configuration file for the virtual machine [Non-Sequential = ID]
CPU_Limit	The cap on the CPU consumption of CPU time by the virtual machine, measured in megahertz (MHz). A value of zero indicates no limit on CPU consumption. [Non-Sequential = SUM]
CPU_Max	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = SUM]
CPU_Min	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = SUM]
CPU_Reservation	The number of CPU cycles reserved for the virtual machine, measured in megahertz (MHz) [Non-Sequential = SUM]
CPU_Shares	The CPU share allocation for the virtual machine [Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Non-Sequential = ID]

Disk_Shares	The disk share allocation for the virtual machine [Non-Sequential = SUM]
ESX_Server	The version of VMware ESX Server [Non-Sequential = ID]
htSharing	Specifies how the VCPUs of a virtual machine are allowed to share physical cores on a hyperthreaded system. Values can be any, internal, or none. [Non-Sequential = ID]
Memory_Limit	The cap on the memory consumption by this virtual machine, measured in megabytes. A value of zero indicates no fixed limit on memory consumption. [Non-Sequential = SUM]
Memory_Max	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = SUM]
Memory_Min	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = SUM]
Memory_Reservation	The amount of memory reserved for the virtual machine, measured in megabytes [Non-Sequential = SUM]
Memory_Shares	The memory share allocation for the virtual machine [Non-Sequential = SUM]
OS	The virtual machine operating system name [Non-Sequential = ID]
PID	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Non-Sequential = ID]
Time	The timestamp of the data sample [Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Non-Sequential = ID]
VM	The virtual machine name [Non-Sequential = ID]
VMID	The virtual machine identifier [Non-Sequential = ID]
VCPU_Count	The number of virtual processors for the virtual machine [Non-Sequential = SUM]

17.2.30. Virtual Machine CPU Usage Table

The Virtual Machine CPU Usage table is derived from the CPU.by Virtual Machine and the VMware.Host Configuration tables. You can view the virtual machine CPU usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Machine CPU Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Performance
Physical tables used to produce this table:	CPU.by Virtual Machine VMware.Host Configuration

Statistic Name	Description
%busy	The percentage of the server processor or processors that the virtual machine used [Sequential = AVG Non-Sequential = SUM]
%vcpu_busy	The percentage of the virtual machines virtual processors used [Sequential = AVG Non-Sequential = SUM]
%vcpu_ready	The percentage of time the virtual machine was ready to perform an operation but had to wait for a processor [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
emin	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = MIN Non-Sequential = SUM]

extrasec	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
guaranteed	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = LST Non-Sequential = SUM]
host_uptime	The elapsed time in seconds between two samples that the host or virtual machine was powered on [Sequential = SUM Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
max	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = MAX Non-Sequential = SUM]
min	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = MIN Non-Sequential = SUM]
ready	The amount of time in seconds the virtual machine was ready to perform an operation but had to wait for a processor [Sequential = SUM Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shares	The number of CPU shares allocated to the virtual machine [Sequential = LST Non-Sequential = SUM]
syssec	The amount of system time in seconds consumed by the virtual machine [Sequential = SUM Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
uptime	The elapsed time in seconds between two samples that the host or virtual machine was powered on [Sequential = SUM Non-Sequential = SUM]
usage(MHz)	The CPU usage in megahertz (MHz) over the collected interval [Sequential = AVG Non-Sequential = SUM]
usedsec	The processor time in seconds consumed by the virtual machine [Sequential = SUM Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
waitsec	The virtual CPU wait time in seconds [Sequential = SUM Non-Sequential = SUM]

17.2.31. Virtual Machine Memory Usage Table

The Virtual Machine Memory Usage table is derived from the Memory.by Virtual Machine and the VMware.Host Configuration tables. You can view the virtual machine memory usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Machine Memory Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Performance
Physical tables used to produce this table:	Memory.by Virtual Machine VMware.Host Configuration

Statistic Name	Description
%usage	The percentage of total available memory that is used [Sequential = AVG Non-Sequential = SUM]
active	The working set size estimate in megabytes at the end of the interval for the virtual machine [Sequential = AVG Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
consumed	The amount of host memory in megabytes consumed by the virtual machine for guest memory [Sequential = AVG Non-Sequential = SUM]
cptread	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = LST Non-Sequential = SUM]
cpttgt	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = LST Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]

DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
max	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = MAX Non-Sequential = SUM]
memctl	The amount of memory in megabytes currently reclaimed using vmmemctl for the virtual machine [Sequential = AVG Non-Sequential = SUM]
memctlgt	The target memory size in megabytes to reclaim using vmmemctl for the virtual machine [Sequential = AVG Non-Sequential = SUM]
min	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = MIN Non-Sequential = SUM]
overhd	The amount of extra memory the virtual machine process is using, in addition to the amount of memory allocated to it in megabytes [Sequential = LST Non-Sequential = SUM]
ovhdmax	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = MAX Non-Sequential = SUM]
ovhdpeak	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = MAX Non-Sequential = SUM]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
shared	The amount of memory in megabytes shared between all running virtual machines and within a virtual machine in megabytes [Sequential = AVG Non-Sequential = SUM]
shares	The number of memory shares allocated to the virtual machine [Sequential = LST Non-Sequential = SUM]
size	The amount of memory in megabytes currently allocated to the virtual machine [Sequential = AVG Non-Sequential = SUM]
sizetgt	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = AVG Non-Sequential = SUM]
swapiin	The total amount of memory in megabytes that has been read from the virtual machine's swap file to the machine memory by the VMKernel during the interval [Sequential = AVG Non-Sequential = SUM]

VMware Systems

swapout	The total amount of memory in megabytes that has been transferred from the virtual machine's swap file to the machine memory by the VMKernel during the interval [Sequential = AVG Non-Sequential = SUM]
swapped	The amount of memory in megabytes currently swapped to the VMware File System 3 (VMFS3) swap file [Sequential = LST Non-Sequential = SUM]
swaptgt	The target size in megabytes to swap to the VMware File System 3 (VMFS3) swap file for the virtual machine [Sequential = LST Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
zero	The amount of memory in megabytes that is zeroed out [Sequential = AVG Non-Sequential = SUM]

17.2.32. Virtual Machine Network Device Usage Table

The Virtual Machine Network Device Usage table is derived from the Network Device.by Virtual Machine and the VMware.Host Configuration tables. You can view the virtual machine network device usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Machine Network Device Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Performance
Physical tables used to produce this table:	Network Device.by Virtual Machine VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KBRx/s	The amount of data received per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KBTx/s	The amount of data transmitted per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KBx/s	The amount of data transferred per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
packets/s	The number of packets transferred per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]

pktsRx/s	The number of packets received per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
pktsTx/s	The number of packets transmitted per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the host system interface [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

17.2.33. Virtual Machine Virtual CPU Usage Table

The Virtual Machine Virtual CPU Usage table is derived from the CPU.by Virtual CPU and the VMware.Host Configuration tables. You can view the virtual machine virtual CPU usage data by cluster, datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Machine Virtual CPU Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Performance
Physical tables used to produce this table:	Network Device.by Virtual Machine VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
extrasec	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = SUM Non-Sequential = SUM]
guaranteed	This statistic is not available for the VMware Infrastructure Agent. The value is reported as <N/A>. [Sequential = LST Non-Sequential = SUM]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
Object	The name of the CPU object [Sequential = ID Non-Sequential = ID]
ready	The virtual CPU time that is spent in the ready state in seconds [Sequential = SUM Non-Sequential = SUM]

Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
syssec	The virtual CPU time that is spent on system processes in seconds [Sequential = SUM Non-Sequential = SUM]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
usage(MHz)	The CPU usage in megahertz (MHz) over the collection interval [Sequential = AVG Non-Sequential = SUM]
usedsec	The virtual CPU time that is used in seconds [Sequential = SUM Non-Sequential = SUM]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]
waitsec	The virtual CPU wait time in seconds [Sequential = SUM Non-Sequential = SUM]

17.2.34. Virtual Machine Network vmnic Usage Table

The Virtual Machine Network vmnic Usage table is derived from the Network Device.vmnic by Virtual Machine and the VMware.Host Configuration tables. You can view the virtual machine network vmnic usage data by cluster, by datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	Virtual Machine Network vmnic Usage
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Virtual Machines/ <i>virtualmachinename</i>
Table type:	Performance
Physical tables used to produce this table:	Network Device.vmnic by Virtual Machine VMware.Host Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Non-Sequential = AVG]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
Datacenter	The name of the datacenter to which the host belongs [Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Host	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters is truncated. [Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = AVG]
KBRx/s	The amount of data received per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
KBTx/s	The amount of data transmitted per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]

KBx/s	The amount of data transferred per second in kilobytes (KB) by the virtual machine [Sequential = AVG Non-Sequential = SUM]
packets/s	The number of packets transferred per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
pktsRx/s	The number of packets received per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
pktsTx/s	The number of packets transmitted per second by the virtual machine [Sequential = AVG Non-Sequential = SUM]
Resource	The name of the host system interface [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the host system. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]
Virtual_Machine	The name of the virtual machine to which the data applies [Sequential = ID Non-Sequential = ID]

17.2.35. VMware CPU Relative Performance Table

The VMware CPU Relative Performance table is derived from the CPU.Relative Performance and the VMware.Storage Configuration tables. You can view the CPU relative performance by cluster, by datacenter, or by VMware vCenter server. This table is only available for viewing in TeamQuest Analyzer.

Table Field Hierarchy

Class:	VMware CPU Relative Performance
Subclass:	
IT Resource Name:	/TeamQuest/System/VMware/Host/ <i>systemname</i>
Table type:	Performance
Physical tables used to produce this table:	CPU.Relative Performance VMware.Storage Configuration

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Cluster	The name of the cluster to which the host belongs. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
ClusterId	The unique identifier for a cluster provided by the VMware API. This identifier does not change when a cluster is renamed. This field contains <N/A> if the host does not belong to a VMware cluster. [Sequential = ID Non-Sequential = ID]
cpu_relative_performance	The relative performance of the CPU on a common scale [Sequential = AVG Non-Sequential = SUM]
Datacenter	The name of the datacenter to which the host belongs [Sequential = ID Non-Sequential = ID]
DatacenterId	The unique identifier for a datacenter provided by the VMware API. This identifier does not change when a datacenter is renamed. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
relative_unused	The amount of CPU resources not used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]
rel_used	The amount of CPU resources used based on a common, relative scale [Sequential = AVG Non-Sequential = SUM]

VMware Systems

System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
vCenter	The name of the VMware vCenter server [Sequential = ID Non-Sequential = ID]

Section 18

Web Server

The Web Server Agent (**tqwsp**) gathers performance information about Web server instances running on your system and stores the information in the aggregation sets of the TeamQuest performance database. It also maintains lists of the top files accessed in the Web File Access table of the performance database.

Statistics are collected on overall connection and error rates, throughput, maximum and average transfer sizes, and average request times. Response codes per second for each category of response code are also stored. In addition, statistics are collected by transfer size, file type, and request type.

This section contains a listing of the statistics collected by the agent:

- Web Server Statistics (see 18.1)
- Top File Access Statistics (see 18.2)

Note: *At the end of each statistic description, you will see a notation in brackets indicating the method that is used for data consolidation (for example, [Sequential = SUM Non-Sequential = SUM]). Sequential means that the field is consolidated over time. Non-Sequential means that the field is consolidated within a specified time interval.*

The following notations are used:

AVG = Average
DIV = Weight
FST = First
ID = Identifier
LST = Last
MAX = Maximum
MIN = Minimum
NON = None or no method was used
SUM = Summation

If you are using TeamQuest View to view aggregation set data, the sequential method is used for data consolidation.

Because derived statistics are not stored in the performance database, the data consolidation method is not shown in the description of a derived statistic.

18.1. Web Server Statistics

The following statistics are stored in the TeamQuest performance database by the Web Server Agent (**tqwsp**). The Resource portion of each parameter name is the *webservername*. The *webservernames* used are those that have been defined in the Web Server Agent configuration file. For more information on configuring the Web Server Agent, see the *TeamQuest Performance Software Administration Guide*.

Note: *The following statistics are only available for the TeamQuest database architecture. If the open database architecture is used, a record for each agent using these statistics is created in the TQ.Agent Interval table.*

Class:	TQ
Subclass:	N/A
IT Resource Name:	N/A
TeamQuest Table Name:	N/A
Open Table Name:	N/A
Statistic Name:	
tqwsp_end_time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
tqwsp_interval	The number of seconds elapsed between the end of data collection for the previous sample and the end of data collection for the current sample [Sequential = SUM Non-Sequential = ID]
wsp interval	The number of seconds elapsed between two samples of the Web Server Agent [Sequential = SUM Non-Sequential = ID]

Table Field Hierarchy

Class:	TQ
Subclass:	Agent Interval
IT Resource Name:	/TeamQuest/System/ <i>systemname</i>
TeamQuest Table Name:	TQ.Agent Interval
Open Table Name:	AGENTINTERVAL
Collection interval:	Based on the collection period
Default retentions:	8 hours at collection period interval 8 days at 10-minute intervals 35 days at 1-hour intervals 400 days at 8-hour intervals
Table type:	Performance
Derived tables using fields from this table:	N/A

Statistic Name	Description
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
Agent	The name of the agent that is collecting data. This field is limited to 52 characters. Any agent name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Instance	The instance name of the agent that is collecting data. This field is limited to 52 characters. Any instance name longer than 52 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
PID	The process identifier of the agent instance that is collecting data [Sequential = ID Non-Sequential = ID]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
System	The name of the system where the data is collected. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential LST Non-Sequential = ID]

Web Server

Parameter Hierarchy

Class:	Web Server
Subclass:	File Types.Bitmap File Types.ColdFusion File Types.Compressed File Types.Document File Types.Dynamic File Types.HTML File Types.Image File Types.Java File Types.Other File Types.PHP File Types.Sound File Types.Video
IT Resource Name:	/TeamQuest/System/ <i>systemname</i> /Web Server/ <i>webservername</i>
TeamQuest Table Name:	Web Server.File Types.Bitmap Web Server.File Types.ColdFusion Web Server.File Types.Compressed Web Server.File Types.Document Web Server.File Types.Dynamic Web Server.File Types.HTML Web Server.File Types.Image Web Server.File Types.Java Web Server.File Types.Other Web Server.File Types.PHP Web Server.File Types.Sound Web Server.File Types.Video
Open Table Name:	WSVRFILETYPESBITMAP WSVRFILETYPESCOLDFUS WSVRFILETYPESCOMPRES WSVRFILETYPESDOCUMENT WSVRFILETYPESDYN WSVRFILETYPESHTML WSVRFILETYPESIMAGE WSVRFILETYPESJAVA WSVRFILETYPESOTHER WSVRFILETYPESPHP WSVRFILETYPESSOUND WSVRFILETYPESVIDEO
Resource:	webserver1, webserver2, ...
Statistic Name:	
% of bytes	The percentage of the total bytes transferred on behalf of connections that requested files of this type [Sequential = AVG Non-Sequential = AVG] View Report: /report/webserv/filetype/bytedist.rpt

% of conn	<p>The percentage of the total connections requesting files of this type</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/webserv/filetype/conndist.rpt</p>
avg req time (secs)	<p>The average number of seconds required to process requests for files of this type (or zero if not available in access log)</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/webserv/filetype/reqtime.rpt</p>
avg xfer size (kbytes)	<p>The average transfer size for files of this type</p> <p>[Sequential = AVG Non-Sequential = AVG]</p> <p>View Report:</p> <p>/report/webserv/filetype/xfersize.rpt</p>
max xfer size (kbytes)	<p>The largest transfer size for files of this type</p> <p>[Sequential = MAX Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/webserv/filetype/xfersize.rpt</p>
xfers/sec	<p>The number of transfers per second for files of this type</p> <p>[Sequential = AVG Non-Sequential = SUM]</p> <p>View Report:</p> <p>/report/webserv/filetype/xferrate.rpt</p>

Web Server

Class:	Web Server
Subclass:	Request Types.Delete Request Types.Get Request Types.Head Request Types.Option Request Types.Other Request Types.Post Request Types.Put Request Types.Trace
IT Resource Name:	/TeamQuest/System/systemname/Web Server/webservername
TeamQuest Table Name:	Web Server.Request Types.Delete Web Server.Request Types.Get Web Server.Request Types.Head Web Server.Request Types.Option Web Server.Request Types.Other Web Server.Request Types.Post Web Server.Request Types.Put Web Server.Request Types.Trace
Open Table Name:	WSREQTYPESDELETE WSREQTYPESGET WSREQTYPESHEAD WSREQTYPESEOPTION WSREQTYPESOTHER WSREQTYPESPOST WSREQTYPESPUT WSREQTYPESTRACE
Resource:	webserver1, webserver2, ...
Statistic Name:	
% of conn	The percentage of the total connections using this HTTP request type [Sequential = AVG Non-Sequential = AVG] View Report: /report/webserv/reqtypes.rpt
avg req time (secs)	The average number of seconds required to process requests using this HTTP request type [Sequential = AVG Non-Sequential = AVG] View Report: /report/webserv/reqtime.rpt
reqs/sec	The number of connections per second using this HTTP request type [Sequential = AVG Non-Sequential = SUM] View Report: /report/webserv/reqrates.rpt

Class: Web Server

Subclass: Response Codes.Client error
Response Codes.Information
Response Codes.Redirection
Response Codes.Server error
Response Codes.Success

IT Resource Name: /TeamQuest/System/systemname/Web Server/webservername

TeamQuest Table Name: Web Server.Response Codes.Client error
Web Server.Response Codes.Information
Web Server.Response Codes.Redirection
Web Server.Response Codes.Server error
Web Server.Response Codes.Success

Open Table Name: WSRSPCODESLIERROR
WSRSPCODESINFORMATIO
WSRSPCODESREDIRECTIO
WSRSPCODESSERVERERROR
WSRSPCODESSUCCESS

Resource: webserver1, webserver2, ...

Statistic Name:
codes/sec The number of this type of HTTP response codes sent per second
[Sequential = AVG Non-Sequential = SUM]
View Report:
/report/webserv/err-rate.rpt

Class: Web Server

Subclass: Summary

IT Resource Name: /TeamQuest/System/systemname/Web Server/webservername

TeamQuest Table Name: Web Server.Summary

Open Table Name: WSVRSUM

Resource: webserver1, webserver2, ...

Statistic Name:
active sessions The number of unique IP addresses that were active during the interval
[Sequential = AVG Non-Sequential = SUM]
avg req time (secs) The average number of seconds required to process an HTTP request (or zero if not available in access log)
[Sequential = AVG Non-Sequential = AVG]
View Report:
/report/webserv/websumm.rpt
avg xfer size (kbytes) The average transfer size for files of this type
[Sequential = AVG Non-Sequential = AVG]
View Report:
/report/webserv/websumm.rpt

Web Server

connections/sec	The number of HTTP requests per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/webserv/webrates.rpt /report/webserv/websumm.rpt
errors/sec	The number of HTTP client or server errors per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/webserv/webrates.rpt /report/webserv/websumm.rpt
max xfer size (kbytes)	The size of the largest transfer in kilobytes [Sequential = MAX Non-Sequential = SUM] View Report: /report/webserv/websumm.rpt
throughput (kbytes/sec)	The number of kilobytes transferred per second [Sequential = AVG Non-Sequential = SUM] View Reports: /report/webserv/webrates.rpt /report/webserv/websumm.rpt
visit count	The number of unique IP addresses whose sessions timed out during the interval [Sequential = SUM Non-Sequential = SUM]

Class:	Web Server
Subclass:	Transfer Sizes.to4KB Transfer Sizes.4to16KB Transfer Sizes.16to64KB Transfer Sizes.64to256KB Transfer Sizes.ov256KB
IT Resource Name:	/TeamQuest/System/systemname/Web Server/webservername
TeamQuest Table Name:	Web Server.Transfer Sizes.to4KB Web Server.Transfer Sizes.4to16KB Web Server.Transfer Sizes.16to64KB Web Server.Transfer Sizes.64to256KB Web Server.Transfer Sizes.ov256KB
Open Table Name:	WSTRSZSTO4KB WSTRSZS4TO16KB WSTRSZS16TO64KB WSTRSZS64TO256KB WSTRSZSOV256KB
Resource:	webserver1, webserver2, ...
Statistic Name:	
% of bytes	The percentage of the total bytes transferred on behalf of connections that requested files of this size [Sequential = AVG Non-Sequential = AVG] View Report: /report/webserv/xfersize/bytedist.rpt
% of conn	The percentage of the total connections requesting files of this size [Sequential = AVG Non-Sequential = AVG] View Report: /report/webserv/xfersize/conndist.rpt
avg req time (secs)	The average number of seconds required to process requests for files of this size (or zero if not available in access log) [Sequential = AVG Non-Sequential = AVG] View Report: /report/webserv/xfersize/reqtime.rpt
xfers/sec	The number of transfers per second for files of this size [Sequential = AVG Non-Sequential = SUM] View Report: /report/webserv/xfersize/xferrate.rpt

18.2. Top File Access Statistics

The Web Server Agent maintains lists of the files being accessed the most on your Web server. The TopN file lists are stored in the TeamQuest performance database tables. A set of file access records is written to the table at each sample period. The maximum number of records to be written at each sample is specified in the Web Server Agent configuration file. The default Web Server Agent configuration file setting saves information on the top 100 files accessed during each sample. If there is no activity on the Web server during a sample period, no records are written.

Table Field Hierarchy

Class:	Web
Subclass:	File Accesses
IT Resource Name:	/TeamQuest/System/systemname/Web Server/webservername
TeamQuest Table Name:	Web.File Accesses
Open Table Name:	WEBFILEACCESSES
Collection interval:	Based on the primary aggregation set
Default retention:	1 day at the primary collection interval 1 month at the secondary 1-hour interval
Table type:	Performance

Statistic Name	Description
accesses	The number of connections that requested the file during the sample period [Sequential = SUM Non-Sequential = SUM]
Actual_Interval	The elapsed time between two samples in seconds. This value may not be the same as the Interval statistic value in all samples because data collection can sometimes take longer than expected or because the associated database became active within the given sample interval. [Sequential = SUM Non-Sequential = ID]
avg_reqtime	The average request time (in seconds) for accesses to the file during the sample period. The value is zero if request time is not available in the access log. [Sequential = AVG Non-Sequential = AVG]
file	The name of the file [Sequential = ID Non-Sequential = ID]
Interval	The expected sampling interval in seconds [Sequential = SUM Non-Sequential = ID]
kbytes	The maximum number of kilobytes transferred for a single access to the file during the sample period [Sequential = MAX Non-Sequential = MAX]

max_reqtime	The longest request time (in seconds) for accesses to the file during the sample period. The value is zero if request time is not available in the access log. [Sequential = MAX Non-Sequential = MAX]
min_reqtime	The shortest request time (in seconds) for accesses to the file during the sample period. The value is zero if request time is not available in the access log. [Sequential = MIN Non-Sequential = MIN]
Sample_End_Time	The timestamp of the actual end of data collection for the current sample [Sequential = LST Non-Sequential = ID]
stddev_reqtime	The standard deviation of the request time (in seconds) for accesses to the file during the sample period. The value is zero if request time is not available in the access log. [Sequential = NON Non-Sequential = NON]
sumsqrs_reqtime	The sum of the squares of the request time (in seconds) for accesses to the file during the sample period. The value is zero if request time is not available in the access log. This is a hidden statistic and is for internal use only. [Sequential = SUM Non-Sequential = SUM]
System	The name of the system on which the Web server resides. This field is limited to 51 characters. Any system name longer than 51 characters will be truncated. [Sequential = ID Non-Sequential = ID]
Time	The timestamp of the data sample [Sequential = LST Non-Sequential = ID]
type	The file type (for example, HTML, Image, Sound, Video, Dynamic, Document, Compressed, Bitmap, Java, or Other). Up to 16 characters are displayed. [Sequential = ID Non-Sequential = ID]
webserver	The Web server name as specified in the Web Server Agent configuration file. Up to 52 characters are displayed. [Sequential = ID Non-Sequential = ID]

Bibliography

- OpenSSL FIPS 140-2 Security Policy, Version 2.0.5*
([OpenSSL-SecurityPolicy-2.0.5.pdf](#)). Open Source Software Institute.
- TeamQuest Administration Console User Guide* (TQ-40321). TeamQuest Corporation.
- TeamQuest AutoPredict Agent Tutorial* (TQ-01426). TeamQuest Corporation.
- TeamQuest CMIS Installation Guide* (TQ-50015). TeamQuest Corporation.
- TeamQuest Performance Software Administration Guide* (TQ-40020).
TeamQuest Corporation.
- TeamQuest Performance Software Administration Reference Manual* (TQ-40022).
TeamQuest Corporation.
- TeamQuest Performance Software Command Line Interfaces Reference Manual*
(TQ-40024). TeamQuest Corporation.
- TeamQuest Performance Software Enterprise Solutions Administration Guide*
(TQ-40212). TeamQuest Corporation.
- TeamQuest Performance Software Installation and Configuration Guide for KVM
Systems* (TQ-17015). TeamQuest Corporation.
- TeamQuest Performance Software Installation and Configuration Guide for
VMware vSphere* (TQ-18015). TeamQuest Corporation.

Index

A

- active SQL statistics, Sybase ASE 15–10
- alarm information, Oracle instances 12–9
- Alarm Service, system alarm
 - statistics 16–1
- alert log alarm statistics, Oracle 12–3
- application statistics, DB2 UDB 3–3
- AutoPredict Agent statistics 2–1

B

- block contention wait statistics,
 - Oracle 12–4
- bufferpool statistics, DB2 UDB 3–17

C

- client statistics, network applications 10–7
- collection agents
 - DB2 UDB Agent 3–1
 - Disk Space Agent
 - HP-UX 4–21
 - IBM AIX 6–34
 - Linux 8–18
 - Oracle Solaris 13–28
 - General Log Agent
 - HP-UX 4–53
 - IBM AIX 6–75
 - Linux 8–50
 - Microsoft Windows 9–40
 - Oracle Solaris 13–69
 - Hyper-V Agent 5–1
 - KVM systems 7–1
 - Network Agent
 - HP-UX 4–25
 - IBM AIX 6–38
 - Linux 8–22
 - Oracle Solaris 13–32

- Network Application Agent 10–1
- Network Device Agent 11–1
- Oracle Alarm Agent 12–1
- Oracle Data Agent 12–1
- PostgreSQL Database Agent 14–1
- Process-Workload Agent
 - HP-UX 4–40
 - IBM AIX 6–62
 - Linux 8–35
 - Microsoft Windows 9–25
 - Oracle Solaris 13–52
- Sybase ASE Agent 15–1
- System Activity Agent
 - HP-UX 4–2
 - IBM AIX 6–3
 - Linux 8–3
 - Microsoft Windows 9–2
 - Oracle Solaris 13–3
- System Log Agent
 - HP-UX 4–52
 - IBM AIX 6–74
 - Linux 8–49
 - Oracle Solaris 13–68
- VMware servers 17–1
- Web Server Agent 18–1
- Windows Event Log Agent 9–42
- Windows Services Agent 9–43
- components of response statistics 2–2
- configuration statistics, Sybase ASE Agent 15–12

D

- data information, Oracle instances 12–10
- database configuration statistics,
 - DB2 UDB 3–23
- database detail statistics
 - PostgreSQL Database Agent 14–3
- database detail statistics,
 - Sybase ASE 15–14
- database statistics, DB2 UDB 3–24

- database status statistics, DB2 UDB 3–41
- database summary statistics
 - PostgreSQL Database Agent 14–2
 - Sybase ASE 15–15
- datafile capacity statistics, Oracle 12–5
- datafile I/O statistics, Oracle 12–7
- DB2 UDB Agent, agent description 3–1
- DB2 UDB Database Server
 - application detail statistics 3–3
 - bufferpool detail statistics 3–17
 - database configuration statistics 3–23
 - database statistics 3–24
 - database status statistics 3–41
 - instance configuration statistics 3–44
 - instance statistics 3–45
 - instance status statistics 3–50
 - table statistics 3–52
 - tablespace statistics 3–54
- derived statistics
 - HP-UX 4–55
 - IBM AIX 6–77
 - Linux 8–52
 - Microsoft Windows 9–45
 - Oracle Solaris 13–71
 - Rules report for TeamQuest View 13–76
- TeamQuest Alert
 - HP-UX 4–59
 - IBM AIX 6–80
 - Linux 8–55
 - Microsoft Windows 9–45
 - Oracle Solaris 13–75
- TeamQuest On the Web
 - HP-UX 4–57
 - IBM AIX 6–79
 - Linux 8–54
 - Microsoft Windows 9–46
 - Oracle Solaris 13–73
- TeamQuest View, Microsoft Windows 9–47
- device detail statistics, Sybase ASE 15–16
- disabling reduction processing
 - HP-UX 4–40
 - IBM AIX 6–62
 - Linux 8–35
 - Microsoft Windows 9–25
 - Oracle Solaris 13–52
- Disk Space Agent
 - HP-UX 4–21
 - IBM AIX 6–34
 - Linux 8–18
 - Oracle Solaris 13–28

- disk space statistics
 - HP-UX 4–21
 - IBM AIX 6–34
 - Linux 8–18
 - Oracle Solaris 13–28

E

- engine detail statistics, Sybase ASE 15–17

G

- General Log Agent
 - HP-UX 4–53
 - IBM AIX 6–75
 - Linux 8–50
 - Microsoft Windows 9–40
 - Oracle Solaris 13–69
- general log statistics
 - HP-UX 4–53
 - IBM AIX 6–75
 - Linux 8–50
 - Microsoft Windows 9–40
 - Oracle Solaris 13–69

H

- hardware configuration information
 - retrieving with **twarp**
 - HP-UX 4–41
 - IBM AIX 6–63
 - Linux 8–35, 8–36
- hardware inventory statistics
 - HP-UX 4–46
 - IBM AIX 6–68
 - Linux 8–43
 - Microsoft Windows 9–33
 - Oracle Solaris 13–60
- host statistics, Hyper-V Agent 5–36
- HP-UX systems
 - derived statistics 4–55
 - disk space statistics 4–21
 - general log statistics 4–53
 - hardware inventory statistics 4–46
 - network statistics 4–25
 - process statistics 4–40
 - system activity statistics 4–2
 - system log statistics 4–52

- TeamQuest Alert statistics 4-59
- TeamQuest log statistics 4-54
- TeamQuest On the Web statistics 4-57
- workload performance derived
 - statistics 4-55
- workload statistics 4-35
- hypervisor statistics, Hyper-V Agent 5-2
- Hyper-V Agent
 - agent description 5-1
 - host statistics 5-36
 - hypervisor statistics 5-2
 - I/O statistics 5-53
 - network statistics 5-56
 - virtual machine statistics 5-24

I

- IBM AIX systems
 - derived statistics 6-77
 - disk space statistics 6-34
 - general log statistics 6-75
 - hardware inventory statistics 6-68
 - LPAR configuration statistics 6-57
 - network statistics 6-38
 - process statistics 6-62
 - system activity 6-3
 - system log statistics 6-74
 - TeamQuest Alert statistics 6-80
 - TeamQuest log statistics 6-76
 - TeamQuest On the Web statistics 6-79
 - workload manager statistics 6-51
 - workload performance derived
 - statistics 6-77
 - workload statistics 6-52
- incomplete process data
 - HP-UX 4-41
 - IBM AIX 6-62
 - Linux 8-35
 - Oracle Solaris 13-52
- instance alarm statistics, Oracle 12-9
- instance configuration statistics,
 - DB2 UDB 3-44
- instance data statistics, Oracle 12-10
- instance statistics, DB2 UDB 3-45
- instance status statistics, DB2 UDB 3-50
- interface statistics, network devices 11-5
- I/O statistics, Hyper-V Agent 5-53

K

- KVM systems, description 7-1

L

- latch statistics, Oracle 12-37
- library cache statistics, Oracle 12-39
- Libvirt Agent
 - agent description 7-1
 - statistics 7-1
- Linux systems
 - derived statistics 8-52
 - disk space statistics 8-18
 - general log statistics 8-50
 - hardware inventory data 8-43
 - LPAR configuration statistics 8-30
 - network statistics 8-22
 - process statistics 8-35
 - system activity statistics 8-3
 - system log statistics 8-49
 - TeamQuest Alert statistics 8-55
 - TeamQuest On the Web statistics 8-54
 - workload performance derived
 - statistics 8-52
 - workload statistics 8-32
- listener alarm statistics, Oracle 12-41
- lock alarm statistics, Oracle 12-42
- lock detail statistics, Sybase ASE 15-19
- LPAR configuration statistics
 - IBM AIX 6-57
 - Linux 8-30

M

- Microsoft Exchange 2000, optional
 - statistics 9-49
- Microsoft Exchange 5.5, optional
 - statistics 9-52
- Microsoft IIS, optional statistics 9-54
- Microsoft SQL Server, optional
 - statistics 9-57
- Microsoft Windows systems
 - derived statistics 9-45
 - general log statistics 9-40
 - hardware inventory statistics 9-33
 - optional system activity statistics 9-48
 - process statistics 9-25
 - system activity statistics 9-2

- TeamQuest log statistics 9–41
- Windows event log statistics 9–42
- Windows services statistics 9–43
- workload statistics 9–20

N

- Network Agent
 - HP-UX 4–25
 - IBM AIX 6–38
 - Linux 8–22
 - Oracle Solaris 13–32
- Network Application Agent, agent
 - description 10–1
- Network Applications
 - client statistics 10–7
 - performance statistics 10–2
 - port statistics 10–9
- Network Device Agent, agent
 - description 11–1
- Network Devices
 - interface statistics 11–5
 - summary statistics 11–2
- network statistics
 - HP-UX 4–25
 - Hyper-V Agent 5–56
 - IBM AIX 6–38
 - Linux 8–22
 - Oracle Solaris 13–32

O

- Oracle Alarm Agent, agent
 - description 12–1
- Oracle Data Agent, agent description 12–1
- Oracle Database Server
 - alert log alarm statistics 12–3
 - block contention wait statistics 12–4
 - datafile capacity statistics 12–5
 - datafile I/O statistics 12–7
 - instance alarm statistics 12–9
 - instance data statistics 12–10
 - latch statistics 12–37
 - library cache statistics 12–39
 - listener alarm statistics 12–41
 - lock alarm statistics 12–42
 - rollback segment statistics 12–43
 - row cache statistics 12–45
 - segment alarm statistics 12–47
 - session statistics 12–49

- session wait statistics 12–53
- system parameters statistics 12–55
- system statistics 12–57
- system wait event statistics 12–58
- top SQL cursors statistics 12–60
- Oracle Solaris systems
 - derived statistics 13–71
 - disk space statistics 13–28
 - general log statistics 13–69
 - hardware inventory statistics 13–60
 - network statistics 13–32
 - process statistics 13–52
 - system activity statistics 13–3
 - system log statistics 13–68
 - TeamQuest Alert statistics 13–75
 - TeamQuest log statistics 13–70
 - TeamQuest On the Web statistics 13–73
 - workload performance derived
 - statistics 13–71
 - workload statistics 13–46

P

- performance statistics
 - network applications 10–2
 - Sybase ASE 15–3
- port statistics, network applications 10–9
- PostgreSQL Database Agent
 - agent description 14–1
 - database detail statistics 14–3
 - database summary statistics 14–2
- process accounting
 - determining off state
 - HP-UX 4–41
 - IBM AIX 6–62
 - Linux 8–35
 - Oracle Solaris 13–52
- process statistics
 - HP-UX 4–40
 - IBM AIX 6–62
 - Linux 8–35
 - Microsoft Windows 9–25
 - Oracle Solaris 13–52
 - Sybase ASE 15–21
- process waits statistics, Sybase ASE 15–24
- Process-Workload Agent
 - disabling reduction processing
 - HP-UX 4–40
 - IBM AIX 6–62
 - Linux 8–35

- Microsoft Windows 9–25
- Oracle Solaris 13–52
- hardware inventory statistics
 - HP-UX 4–46
 - IBM AIX 6–68
 - Linux 8–43
 - Microsoft Windows 9–33
 - Oracle Solaris 13–60
- HP-UX 4–40
- IBM AIX 6–62
- incomplete data
 - HP-UX 4–41
 - IBM AIX 6–62
 - Linux 8–35
 - Oracle Solaris 13–52
- Linux 8–35
- Microsoft Windows 9–25
- Oracle Solaris 13–52
- process accounting turned off
 - HP-UX 4–41
 - IBM AIX 6–62
 - Linux 8–35
 - Oracle Solaris 13–52
- reduced process records
 - HP-UX 4–40
 - IBM AIX 6–62
 - Linux 8–35
 - Microsoft Windows 9–25
 - Oracle Solaris 13–52
- retrieving hardware configuration
 - information 13–52
- workload statistics collected by 9–20

R

- reduced process records
 - HP-UX 4–40
 - IBM AIX 6–62
 - Linux 8–35
 - Microsoft Windows 9–25
 - Oracle Solaris 13–52
- retrieving hardware configuration
 - information
 - HP-UX 4–41
 - IBM AIX 6–63
 - Linux 8–35, 8–36
 - Microsoft Windows 9–26
 - Oracle Solaris 13–52

- retrieving hardware configuration
 - information with **tqwarp**
 - HP-UX 4–41
 - IBM AIX 6–63
 - Linux 8–35, 8–36
 - Microsoft Windows 9–26
 - Oracle Solaris 13–52
- rollback segment statistics, Oracle 12–43
- row cache statistics, Oracle 12–45
- Rules report statistics for TeamQuest
 - View 13–76

S

- segment alarm statistics, Oracle 12–47
- session statistics, Oracle 12–49
- session wait statistics, Oracle 12–53
- stretch factor statistics 2–8
- summary statistics, network devices 11–2
- Sybase ASE Agent
 - active SQL statistics 15–10
 - configuration statistics 15–12
 - database detail statistics 15–14
 - database summary statistics 15–15
 - description 15–1
 - device detail statistics 15–16
 - engine detail statistics 15–17
 - lock detail statistics 15–19
 - performance statistics 15–3
 - process statistics 15–21
 - process waits statistics 15–24
 - system waits statistics 15–25
- System Activity Agent
 - HP-UX 4–2
 - IBM AIX 6–3
 - Linux 8–3
 - Microsoft Exchange 2000 statistics 9–49
 - Microsoft Exchange 5.5 statistics 9–52
 - Microsoft IIS statistics 9–54
 - Microsoft SQL Server statistics 9–57
 - Microsoft Windows 9–2
 - optional statistics 9–48
 - Oracle Solaris 13–3
- system activity statistics
 - HP-UX 4–2
 - IBM AIX 6–3
 - Linux 8–3
 - Microsoft Windows 9–2
 - Oracle Solaris 13–3
- System Alarms table statistics 16–1

- System Log Agent
 - HP-UX 4–52
 - IBM AIX 6–74
 - Linux 8–49
 - Oracle Solaris 13–68
- system log statistics
 - HP-UX 4–52
 - IBM AIX 6–74
 - Linux 8–49
 - Oracle Solaris 13–68
- system parameters statistics,
 - Oracle 12–55
- system statistics, Oracle 12–57
- system wait event statistics, Oracle 12–58
- system waits statistics, Sybase ASE 15–25

T

- table classification scheme 1–6
- table statistics, DB2 UDB 3–52
- tablespace statistics, DB2 UDB 3–54
- TeamQuest Alert statistics
 - HP-UX 4–59
 - IBM AIX 6–80
 - Linux 8–55
 - Microsoft Windows 9–45
 - Oracle Solaris 13–75
- TeamQuest log statistics
 - HP-UX 4–54
 - IBM AIX 6–76
 - Microsoft Windows 9–41
 - Oracle Solaris 13–70
- TeamQuest On the Web statistics
 - HP-UX 4–57
 - IBM AIX 6–79
 - Linux 8–54
 - Microsoft Windows 9–46
 - Oracle Solaris 13–73
- TeamQuest View Rules report
 - statistics 13–76
- TeamQuest View, Microsoft Windows
 - derived statistics 9–47
- top SQL cursors statistics, Oracle 12–60
- tqalm**, description 16–1
- tqbnp** agent statistics
 - HP-UX 4–25
 - IBM AIX 6–38
 - Linux 8–22
 - Oracle Solaris 13–32
- tqbbsp** agent statistics
 - HP-UX 4–2, 4–55
 - IBM AIX 6–3, 6–77
 - Linux 8–3, 8–52
 - Microsoft Windows 9–2
 - Oracle Solaris 13–3, 13–71
- tqdb2p** agent statistics
 - agent description 3–1
 - application 3–3
 - bufferpool 3–17
 - database 3–24
 - database configuration 3–23
 - database status 3–41
 - instance 3–45
 - instance configuration 3–44
 - instance status 3–50
 - table 3–52
 - tablespace 3–54
- tqdsp** agent statistics
 - HP-UX 4–21
 - IBM AIX 6–34
 - Linux 8–18
 - Oracle Solaris 13–28
- tqglp** agent statistics
 - HP-UX 4–53
 - IBM AIX 6–75
 - Linux 8–50
 - Microsoft Windows 9–40
 - Oracle Solaris 13–69
- tqhyperv** agent statistics
 - agent description 5–1
 - host 5–36
 - hypervisor 5–2
 - I/O 5–53
 - network 5–56
 - virtual machine 5–24
- tqlibvirt** agent statistics, agent
 - description 7–1
- tqlog** agent statistics
 - HP-UX 4–54
 - IBM AIX 6–76
 - Microsoft Windows 9–41
 - Oracle Solaris 13–70
- tqmdl** agent statistics 2–1
- tqndp** agent statistics
 - agent description 11–1
 - interface 11–5
 - summary 11–2
- tqnpp** agent statistics
 - agent description 10–1
 - client 10–7

- performance 10–2
- port 10–9
- tqoraalm** agent statistics
 - agent description 12–1
 - alert log alarm 12–3
 - instance alarm 12–9
 - instance data 12–10
 - listener alarm 12–41
 - lock alarm 12–42
 - segment alarm 12–47
- tqorap** agent statistics
 - agent description 12–1
 - block contention wait 12–4
 - datafile capacity 12–5
 - datafile I/O 12–7
 - latch 12–37
 - library cache 12–39
 - rollback segment 12–43
 - row cache 12–45
 - session 12–49
 - session wait 12–53
 - system 12–57
 - system parameters 12–55
 - system wait event 12–58
 - top SQL cursors 12–60
- tqpgsql** agent statistics
 - agent description 14–1
 - database detail 14–3
 - database summary 14–2
- tqslp** agent statistics
 - HP-UX 4–52
 - IBM AIX 6–74
 - Linux 8–49
 - Oracle Solaris 13–68
- tqsybase** agent statistics
 - active SQL 15–10
 - agent description 15–1
 - configuration 15–12
 - database detail 15–14
 - database summary 15–15
 - device detail 15–16
 - engine detail 15–17
 - lock detail 15–19
 - performance statistics 15–3
 - process 15–21
 - process waits 15–24
 - system waits 15–25
- tqvmwarep** agent statistics
 - agent description 17–1
- tqwarp** agent statistics
 - disabling reduction processing
 - HP-UX 4–40
 - IBM AIX 6–62
 - Linux 8–35
 - Microsoft Windows 9–25
 - Oracle Solaris 13–52
 - hardware 4–41
 - hardware configuration information
 - HP-UX 4–41
 - IBM AIX 6–63
 - Linux 8–35, 8–36
 - Microsoft Windows 9–26
 - Oracle Solaris 13–52
 - HP-UX
 - hardware inventory statistics 4–46
 - process statistics 4–40
 - workload performance derived
 - statistics 4–55
 - workload statistics 4–35
 - IBM AIX
 - hardware inventory statistics 6–68
 - process statistics 6–62
 - workload manager 6–51
 - workload performance derived
 - statistics 6–77
 - workload statistics 6–52
 - Linux
 - hardware inventory statistics 8–43
 - process statistics 8–35
 - workload performance derived
 - statistics 8–52
 - workload statistics 8–32
 - Microsoft Windows 9–25
 - Oracle Solaris
 - hardware inventory statistics 13–60
 - process statistics 13–52
 - workload performance derived
 - statistics 13–71
 - workload statistics 13–46
 - process accounting turned off
 - HP-UX 4–41
 - IBM AIX 6–62
 - Linux 8–35
 - Oracle Solaris 13–52
 - reduced process records
 - HP-UX 4–40
 - IBM AIX 6–62
 - Linux 8–35
 - Microsoft Windows 9–25
 - Oracle Solaris 13–52

Index

- tqwsp** agent statistics
 - agent description 18–1
 - top file access 18–10
 - Web server statistics 18–2
- tqw2kevent**, agent description 9–42
- tqw2kserv**, agent description 9–43
- tqw2kwarp**, agent description 9–25

V

- virtual machine statistics, Hyper-V
 - Agent 5–24
- VMware servers, description 17–1

W

- Web Server Agent
 - agent description 18–1
 - File Accesses table 18–10
 - statistics 18–2
- Windows Event Log Agent
 - agent description 9–42
 - statistics collected by 9–42
- Windows Services Agent
 - agent description 9–43
 - statistics collected by 9–43
- Win32 performance data API
 - limitations 9–3
- workload
 - derived statistics
 - HP-UX 4–55
 - IBM AIX 6–77
 - Linux 8–52
 - Oracle Solaris 13–71
 - statistics
 - HP-UX 4–35
 - IBM AIX 6–52
 - Linux 8–32
 - Microsoft Windows 9–20
 - Oracle Solaris 13–46
- workload manager statistics, IBM
 - AIX 6–51